arXiv:1608.04362v2 [cs.CR] 25 Oct 2016

Computational Soundness for Dalvik Bytecode

Michael Backes
CISPA, Saarland University & MPI-SWS
Saarland Informatics Campus

Robert Künnemann
CISPA, Saarland University
Saarland Informatics Campus

Esfandiar Mohammadi
ETH Zurich

Abstract

Automatically analyzing information flow within Android applications that rely on cryptographic operations with their computational security guarantees imposes formidable challenges that existing approaches for understanding an app's behavior struggle to meet. These approaches do not distinguish cryptographic and non-cryptographic operations, and hence do not account for cryptographic protections: $f(m)$ is considered sensitive for a sensitive message m irrespective of potential secrecy properties offered by a cryptographic operation f. These approaches consequently provide a safe approximation of the app's behavior, but they mistakenly classify a large fraction of apps as potentially insecure and consequently yield overly pessimistic results.

In this paper, we show how cryptographic operations can be faithfully included into existing approaches for automated app analysis. To this end, we first show how cryptographic operations can be expressed as symbolic abstractions within the comprehensive Dalvik bytecode language. These abstractions are accessible to automated analysis and can be conveniently added to existing app analysis tools using minor changes in their semantics. Second, we show that our abstractions are faithful by providing the first computational soundness result for Dalvik bytecode, i.e., the absence of attacks against our symbolically abstracted program entails the absence of any attacks against a suitable cryptographic program realization. We cast our computational soundness result in the CoSP framework, which makes the result modular and composable.

Contents

1 Introduction **4**
- 1.1 Our Contributions . . . 4
- 1.2 Summary of Our Techniques . . . 5
- 1.3 Overview . . . 8

2 Notation **8**

3 CoSP Framework (Review) **8**
- 3.1 Symbolic Model & Execution . . . 9
- 3.2 Symbolic Equivalence . . . 11
- 3.3 Computational execution . . . 13
- 3.4 Computational Soundness . . . 15

4 Dalvik Bytecode (Review) **15**
- 4.1 Syntax of ADL . . . 15
- 4.2 The Semantics of ADL . . . 17

5 Security framework **19**
- 5.1 Execution and communication model . . . 19
- 5.2 Threat model . . . 20
- 5.3 Calls to crypto APIs . . . 21
- 5.4 Indistinguishability of two ADL programs . . . 24

6 Symbolic Dalvik Bytecode **24**
- 6.1 Embeddable symbolic CoSP models . . . 24
- 6.2 Semantics of symbolic ADL . . . 26

7 Computational Soundness **27**
- 7.1 Split-state semantics . . . 28
- 7.2 Split-state representation of ADL . . . 30
- 7.3 Over-approximating ADL . . . 35
- 7.4 Canonical symbolic semantics . . . 41
- 7.5 Constructing the CoSP-embedding . . . 42
  - 7.5.1 Instantiating the over-approximated split-state semantics with references . . . 42
  - 7.5.2 Constructing the algorithm for the recursion step . . . 43
  - 7.5.3 Corecursively defining the embedding . . . 45
- 7.6 Soundness of the CoSP-embedding . . . 46
  - 7.6.1 Preliminary definitions . . . 46
  - 7.6.2 The proofs . . . 47

8 Related work **50**

9 Conclusion and future work **52**

10 Acknowledgements **52**

Appendix **55**

A Inference rules for the DEX Bytecode Semantics **55**

B Extendability **55**
B.1 Arbitrary bit-operations on bitstrings 57
B.2 Derived destructors with symbolic operations 59

1 Introduction

Android constitutes an open-source project not only in terms of source code but also in terms of the whole ecosystem, allowing practically everyone to program new apps and make them publicly available in Google Play. This open nature of Android has facilitated a rapid pace of innovation, but it has also led to the creation and widespread deployment of malicious apps [1, 2]. Such apps often cause privacy violations that leak sensitive information such as location or the user's address book, either as an intended functionality or as a result of uninformed programming. In some cases such apps can even extract sensitive information from honest apps.

A comprehensive line of research has, hence, strived to rigorously analyze how apps are accessing and processing sensitive information. These approaches typically employ the concept of information flow control (IFC), i.e., certain information sources such as GPS position and address book are declared to be sensitive, and certain information sinks are declared to be adversarially observable. An IFC-based analysis then traces the propagation of sensitive information through the program, i.e., if sensitive data m is input to a function f, then the result $f(m)$ is considered sensitive as well. IFC-based analyses thereby determine if information from sensitive sources can ever reach an observable sink, and in that case report a privacy violation.

A considerable number of apps rely on cryptographic operations, e.g., for encrypting sensitive information before it is sent over the Internet. However, analyzing information flow within Android apps that rely on such cryptographic operations with their computational security guarantees imposes formidable challenges that all existing approaches for automated app analysis struggle to meet, e.g., [3–5]. Roughly, these approaches do not distinguish cryptographic operations from other, non-cryptographic functions. Thus, the standard information-tracing mechanism for arbitrary functions applies: $f(m)$ is considered sensitive for a sensitive message m irrespective of potential secrecy properties offered by a cryptographic function f, e.g., the encryption of a sensitive message m is still considered sensitive such that sending this encryption over the Internet is considered a privacy breach. These approaches consequently provide a safe approximation of the app's behavior, but they mistakenly classify a large fraction of apps as potentially insecure and consequently yield *overly pessimistic results*. While approaches based on manual declassification have successfully managed to treat cryptographic operations and their protective properties more accurately, see the section on related work for more details, no concept for an accurate cryptographic treatment is known for automated analysis of Android apps.

1.1 Our Contributions

In this paper, we show how cryptographic operations can be faithfully included into existing approaches for automated app analysis on Android, in the presence of malicious apps or network parties aiming to extract sensitive information from honest parties. Our paper makes two main tangible contributions to this field: (i) we show how cryptographic operations can be expressed as symbolic abstractions within Dalvik bytecode, so that existing automated analysis tools can adopt them with only minor changes in their semantics; and (ii) we show that our abstractions are faithful by providing the first computational soundness result for the comprehensive Dalvik bytecode language, i.e., the absence of attacks against our symbolically abstracted program entails the absence of any attacks against a suitable cryptographic program realization.

Symbolic abstractions in Dalvik bytecode. We first show how cryptographic operations can be expressed as symbolic abstractions within Dalvik bytecode. These symbolic abstractions – often also referred to as perfect cryptography or Dolev-Yao models – constitute idealizations of cryptographic operations as free algebras that lend themselves towards automated analysis. Deriving such abstractions within the comprehensive Dalvik bytecode language constitutes a challenging task, since existing formalizations of Dalvik do not offer a distinction between honest and adversarially controlled components, which is crucial for defining the rules that the symbolic adversary has to adhere to. To

this end, we develop a novel semantic characterization of Dalvik bytecode that we call *split-state semantics* that provides a clear separation between honest program parts and cryptographic API calls with their corresponding augmented adversarial symbolic capabilities. Moreover, this split-state form is key to our proof of computational soundness, see below. Existing tools for automated app analysis can conveniently include our abstractions using minor changes in their underlying semantics, and thereby reason more accurately about cryptographic operations.

Computational soundness for Dalvik bytecode. We show that our symbolic abstractions can be securely instantiated using suitable cryptographic primitives, and thereby provide the first computational soundness result for Dalvik bytecode. More specifically, our result is grounded in the Abstract Dalvik Language (ADL) [4], which constitutes the currently most detailed and comprehensive operational semantics for Dalvik in the literature. To this end, we first extended ADL by probabilistic choices, as it otherwise would be inappropriate to express cryptographic operations.

We cast our computational soundness result in CoSP, a framework for establishing computational soundness results that decouples the process of embedding programming languages into CoSP from the computational soundness proofs itself. In particular, by casting our soundness results in CoSP, a rich body of computational soundness results for individual cryptographic primitves [6–11] is immediately valid for Dalvik bytecode without any additional work.

Establishing computational soundness results for Dalvik bytecode imposed a series of technical challenges that many prior computational soundness works did not have to cope with. We highlight one such challenge. Computational soundness results struggle when confronted with situations in which binary operations are applied to outputs of cryptographic operations, e.g., if the parity of a ciphertext should be checked, since such an operation would be undefined in the symbolic setting. Prior computational soundness results simply excluded programs with such illegitimate operations; this exclusion introduced an additional proof obligation for the automated analysis tool. While excluding such programs simplifies the soundness result, integrating these additional proof obligations in existing automated app analysis tools constitutes a tedious task, since the tools would need to check upfront whether any symbolically undefined operations will be performed on symbolic terms, in any execution branch. As a consequence, we hence decided to establish a computational soundness result that overapproximates such scenarios by sending information of such illegitimate operations to the adversary and letting the adversary decide the result of such operations.

Finally, our proof reveals an additional result that we consider of independent interest: we show that any small-step semantics S in split-state form entails a canonical small-step semantics S^* for a symbolic model that is computationally sound with respect to S. Hence, for establishing a computationally sound symbolic abstraction for *any* given programming language, it suffices to show that the interaction with the attacker and the cryptographic API can be expressed by means of our concept of split-state semantics.

1.2 Summary of Our Techniques

This section summarizes the techniques that we use to achieve these results. We believe that this makes the paper better accessible. Finally, we discuss how our results can be used to extend information flow tools.

The CoSP framework. A central idea of our work is to reduce computational soundness of Dalvik bytecode to computational soundness in the CoSP framework [6, 10]. All definitions in CoSP are cast relative to a *symbolic model* that specifies a set of constructors and destructors that symbolically represent the cryptographic operations and are also used for characterizing the terms that the attacker can derive (called *symbolic attacker knowledge*), and a *computational implementation* that specifies cryptographic algorithms for these constructors and destructors. In CoSP, a *protocol* is represented by an infinite tree that describes the protocol as a labeled transition system. Such a CoSP protocol contains actions for performing abstract computations (applying constructors and destructors to messages) and for communicating with an adversary. A CoSP protocol is equipped with two different

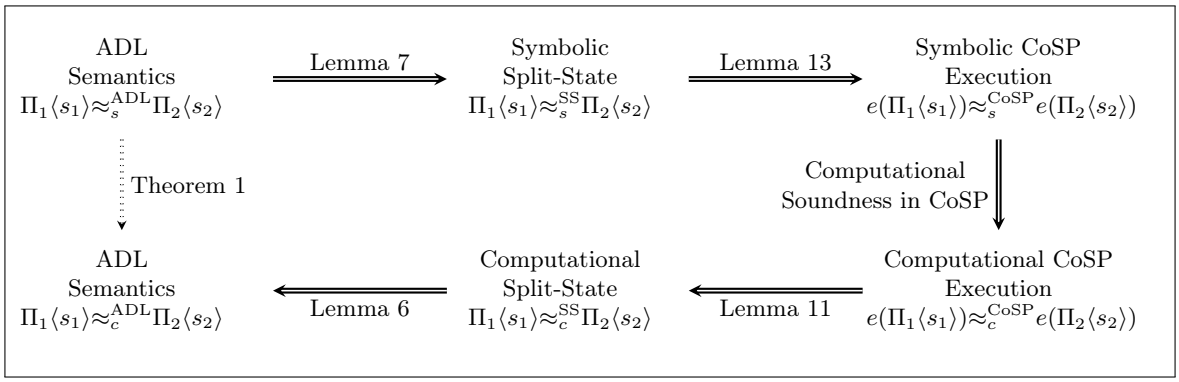

Figure 1: Overview of the main technical lemmas, where e is the embedding into CoSP

semantics: (i) a *symbolic CoSP execution*, in which messages are represented by terms; and (ii) a *computational CoSP execution*, in which messages are bitstrings, and the computational implementation is used instead of applying constructors and destructors. A computational implementation is said to be *computationally sound* for a class of security properties if any CoSP protocol that satisfies these properties in the symbolic execution also satisfies these properties in the computational execution. The advantage of expressing computational soundness results in CoSP is that the protocol model in CoSP is very general so that the semantics of other languages can be embedded therein, thereby transferring the various established soundness results from CoSP into these languages [6–11].

Symbolic ADL and probabilistic choices. ADL as defined in [4] does not support probabilistic choices, and hence no generation of cryptographic keys and no executions of cryptographic functions. We thus extended ADL with a rule that uniformly samples random values (more concretely: a register value from the set of numerical values).

We consider attackers that are external to the app, e.g., malicious parties or network parties with the goal of extracting secrets from an honest app. We characterize the interaction points of the attacker with our extended version of ADL by a set of so-called *malicious functions* that communicate with the attacker. The attacker itself is modeled as a probabilistic polynomial-time machine. We introduce an additional semantic rule that is applied whenever a malicious function is called. The rule invokes the attacker with the arguments of the function call and stores the response of the attacker as a return value. For defining the indistinguishability of two ADL programs, we additionally require that the adversary can also send a single bit b as a final guess, similar to other indistinguishability definitions. This entails the notions of *symbolic equivalence* ($\approx_s^{\mathrm{ADL}}$, in the symbolic setting) and of *computational indistinguishability* ($\approx_c^{\mathrm{ADL}}$, in the computational setting) of two ADL programs.

Split-state semantics for symbolic ADL. Establishing a computational soundness proof for ADL requires a clear separation between honest program parts and cryptographic API calls with their corresponding augmented adversarial symbolic capabilities. To achieve this, we characterize this partitioning by introducing the concept of a *split-state form* of an operational semantics. The split-state form partitions the original semantics into three components, parallelly executed asynchronously: (*i*) all steps that belong to computing cryptographic operations (called the *crypto-API semantics*), (*ii*) all steps that belong to computing the malicious functions (called the *attacker semantics*), and (*iii*) all steps that belong to the rest of the program (called the *honest-program semantics*). Moreover, we define explicit transitions between each of these components, which gives rise to a precise message-passing interface for cryptographic operations and for communicating with the attacker.

Our strategy for showing computational soundness is to use this split-state form for phrasing the symbolic variant as a small-step semantics by replacing the crypto-API semantics with the symbolic constructors and destructors from the symbolic model, and by replacing the attacker semantics by the symbolic characterization of the attacker. With this symbolic semantics at hand, we define split-state

symbolic equivalence ($\approx_s^{\mathrm{SS}}$) as equivalence of (sets of) traces. However, as explained before, we first have to resolve the problem that computational soundness results struggle to deal with situations in which binary operations are applied to outputs of cryptographic operations. We decided not to exclude programs that exhibit such behaviors, but to perform an over-approximation instead by letting the adversary determine the outcome of such operations. This makes our abstractions conveniently accessible for existing tools, but it also complicates the computational soundness proof since we have to consider the operations of constructors and destructors on non-symbolic terms as well. To this end, we encode them as bitstrings and interpret these bitstrings symbolically again. Fortunately, symbolic bitstring interpretations can be seamlessly combined with all previous CoSP results.

We finally define computational indistinguishability of two honest program semantics in the split-state computational execution of ADL ($\approx_c^{\mathrm{SS}}$). As usual, the adversary we consider is a probabilistic polynomial-time machine, and all semantics constitute families of semantics that are indexed by a security parameter.

It might be of independent interest that our symbolic variant of the semantics and the computational indistinguishability can be defined on the split-state form independently of ADL. We show that ADL can be brought into such a split-state form, then prove later that symbolic equivalence in ADL implies symbolic equivalence in the split-state form, and conclude by proving that computational indistinguishability in the split-state form implies computational indistinguishability in ADL. Hence, for every two ADL programs Π_1, Π_2 and initial configurations s_1, s_2 ($\Pi_i\langle s_i\rangle$ denoting Π_i with initial configuration s_i) we have

$$\Pi_1\langle s_1\rangle \approx_s^{\mathrm{ADL}} \Pi_2\langle s_2\rangle \implies \Pi_1\langle s_1\rangle \approx_s^{\mathrm{SS}} \Pi_2\langle s_2\rangle, \text{ and}$$
$$\Pi_1\langle s_1\rangle \approx_c^{\mathrm{SS}} \Pi_2\langle s_2\rangle \implies \Pi_1\langle s_1\rangle \approx_c^{\mathrm{ADL}} \Pi_2\langle s_2\rangle.$$

Computational soundness proof. We first construct an injective embedding e that maps every ADL program to a CoSP protocol. We stress that within CoSP, the same CoSP protocol is used for the computational and the symbolic execution and that CoSP requires a separation of the attacker and the cryptographic operations from the rest of the program. Our split-state form precisely satisfies these requirements. The embedding e uses the honest-program semantics to iteratively construct a CoSP protocol: a transition to the crypto-API semantics corresponds to a computation node; a transition to the attacker semantics corresponds to an output node followed by an input node; and whenever several possibilities exist, a control node is selected to let the adversary decide which possibility (which node) to take.

We prove this embedding sound in the symbolic model, and we prove it complete with respect to the range of e in the computational model, i.e., for every two ADL programs Π_1, Π_2 and initial configurations s_1, s_2 we have

$$\Pi_1\langle s_1\rangle \approx_s^{\mathrm{SS}} \Pi_2\langle s_2\rangle \implies e(\Pi_1\langle s_1\rangle) \approx_s^{\mathrm{CoSP}} e(\Pi_2\langle s_2\rangle)$$

and

$$e(\Pi_1\langle s_1\rangle) \approx_c^{\mathrm{CoSP}} e(\Pi_2\langle s_2\rangle) \implies \Pi_1\langle s_1\rangle \approx_c^{\mathrm{SS}} \Pi_2\langle s_2\rangle,$$

where $\approx_s^{\mathrm{CoSP}}$ and $\approx_c^{\mathrm{CoSP}}$ denote symbolic equivalence and computational indistinguishability in CoSP.

Figure 1 finally shows how all pieces are put together:

Theorem 1 (computational soundness of Dalvik – simplified). *Let Π_1, Π_2 be two ADL programs that use the same crypto-API and s_1, s_2 be two initial configurations. Then we have*

$$\Pi_1\langle s_1\rangle \approx_s^{\mathrm{ADL}} \Pi_2\langle s_2\rangle \implies \Pi_1\langle s_1\rangle \approx_c^{\mathrm{ADL}} \Pi_2\langle s_2\rangle.$$

Extension of information flow tools. To put our work in perspective, we elaborate on a possible application of our result. We envision the extension of information flow (IF) methods with symbolic

abstractions. On a high-level, we envision the following approach for extending IF-tools: whenever there is a potential information flow from High to Low and a cryptographic function is called, we extract a model of the app and query a symbolic prover to find our whether the attacker learns something about the High values. As in symbolic ADL only a few semantic rules are changed w.r.t. ADL, modifications to existing analyses are likely to be confined, thus simple to integrate.

This approach imposes the challenge of extracting a model of the app. Our embedding of an ADL program into CoSP already extracts a symbolic model for the program. However, shrinking this extracted model to a manageable size and querying the symbolic prover in a way such that it scales to complex apps is a task far from simple that merits a paper on its own.

1.3 Overview

Section 3 reviews the CoSP framework for equivalence properties that we ground our computational soundness result on. Section 4 reviews the Abstract Dalvik Language (ADL). Sections 5 and 6 define the probabilistic execution of ADL and introduce our symbolic variant of ADL, including the symbolic abstractions of cryptographic operations and the capabilities of the symbolic adversary. Section 7 defines the connections between ADL, symbolic ADL, and CoSP, and based on these connections proves the computational soundness result. Section 8 discusses related work. We conclude in Section 9 with a summary of our findings and outline directions for future research.

2 Notation

Let $\mathbb{N}$ be the set of natural numbers and assume that they begin at 0. We abbreviate the statement that a set A is a finite subset of a set B as $A \subseteq_{fin} B$. For indicating that function f from a set A to a set B is a partial function, we write $f : A \rightharpoonup B$. We use squared brackets in two different ways: (i) $m[pp]$ denotes the instruction with the number pp for a set of instructions m, and (ii) $r[v \mapsto val] := (r \setminus (v, r(v))) \cup (v, val)$ is short-hand for the function mapping v to val and otherwise behaving like r. We lift this notation to functions, i.e., $r[f](v)$ is equal to $f(v)$ if f is defined on v, and equal to $r(v)$ otherwise. Throughout the paper, we use η as the security parameter. We use ε to denote the empty sequence, empty path, empty bitstring, or empty action depending on the context. We write $\underline{t}$ for a sequence $t_1, \ldots, t_n$ if n is clear from the context. For any sequence $l \in E^*$, we use $\cdot$ to denote concatenation $l_1 \cdot l_2$, as well as the result of appending $(l \cdot e)$ or prepending $(e \cdot l)$ an element, as long as the difference is clear from context. We use $(f(l))_{e \in l}$ to the sequence resulting from applying the (meta-language) operation f to each element in a sequence l, e.g., $(e^2)_{e \in \{1,2,3\}}$ is the sequence of the first three square numbers. We use $l|_0^k$ to denote the k-prefix of l. We filter a sequence l by a set S, denoted $l|_S$, by removing each element that is not in S. We use similar notation for the projection of a sequence $l = s_1, \ldots, s_n \in S^*$: given a partial function $\pi \colon S \rightharpoonup T$, $l|_\pi = (s_1, \ldots, s_{n-1})|_\pi \cdot \pi(s_n)$ or $(s_1, \ldots, s_{n-1})|_\pi$, if $\pi(s_n)$ undefined. $(e)^l$ denotes the sequence of length l where each element equals e. As we represent the attacker as a transition system, we sometimes write T_A and sometimes $\mathcal{A}$ for the attacker.

3 CoSP Framework (Review)

The computational soundness proof developed in this paper follows CoSP [6,10], a general framework for conducting computational soundness proofs of symbolic cryptography and for embedding these proofs into programming languages with their given semantics. CoSP enables proving computational soundness results in a conceptually modular and generic way: every computational soundness result for a cryptographic abstraction phrased in CoSP automatically holds for all suitably embedded languages, and the process of embedding is conceptually decoupled from computational soundness proofs.

Hence in this work, we will suitably embed Dalvik Bytecode into CoSP, and thereby leverage existing computational soundness results of CoSP.

In Section 3.1, we review the *symbolic model* of CoSP, which encompasses symbolic abstractions of cryptographic operations and the representation of programs (which are called protocols) in CoSP. In Section 3.2, we present the notion of *symbolic equivalence* in CoSP, which defines that two CoSP protocols are indistinguishable for an attacker that operates on symbolic abstractions. In Section 3.3, we review the notion of *computational indistinguishability* in CoSP, which defines the execution of protocols using actual cryptographic algorithms. In Section 3.4, we finally review CoSP's notion of *computational soundness*.

3.1 Symbolic Model & Execution

CoSP provides a general symbolic model for expressing cryptographic abstractions. We start with some basic concepts such as constructors, destructors, nonces and message types.

Definition 1 (CoSP Terms)**.** *A* constructor C *is a symbol with a (possibly zero) arity. A* nonce N *is a symbol with zero arity. We write* $C/n \in \mathbf{C}$ *to denote that* $\mathbf{C}$ *contains a constructor* C *with arity* n*. A* message type $\mathbf{T}$ over $\mathbf{C}$ and $\mathbf{N}$ *is a set of terms over constructors* $\mathbf{C}$ *and nonces* $\mathbf{N}$*. A* destructor D *of arity* n*, written* D/n*, over a message type* $\mathbf{T}$ *is a partial map* $\mathbf{T}^n \to \mathbf{T}$*. If* D *is undefined on* $\underline{t}$*, we write* $D(\underline{t}) = \bot$*.*

In CoSP, symbolic abstractions of protocols and of the attacker are formulated in a symbolic model, including countably infinite set of nonces partitioned into protocol and attacker nonces.

Definition 2 (Symbolic model)**.** *A* symbolic model $\mathbf{M} = (\mathbf{C}, \mathbf{N}, \mathbf{T}, \mathbf{D})$ *consists of a set of constructors* $\mathbf{C}$*, a set of nonces* $\mathbf{N}$*, a message type* $\mathbf{T}$ *over* $\mathbf{C}$ *and* $\mathbf{N}$ *(with* $\mathbf{N} \subseteq \mathbf{T}$*), a set of destructors* $\mathbf{D}$ *over* $\mathbf{T}$*. We require that* $\mathbf{N} = \mathbf{N_E} \uplus \mathbf{N_P}$ *for countably infinite sets* $\mathbf{N_P}$ *of protocol nonces and attacker nonces* $\mathbf{N_E}$*.*

To unify notation for constructors, destructors, and nonces, we define the *evaluation of terms* as the partial function $eval_F : \mathbf{T}^n \to \mathbf{T}$ for every constructor or destructor $F/n \in \mathbf{D} \cup \mathbf{C}$ and every nonce $F \in \mathbf{N}$ as follows: (where $n = 0$ for a nonce).

Definition 3 (Evaluation of terms)**.** *If* F *is a constructor, define* $eval_F(\underline{t}) := F(\underline{t})$ *if* $F(\underline{t}) \in \mathbf{T}$ *and* $eval_F(\underline{t}) := \bot$ *otherwise. If* F *is a nonce, define* $eval_F() := F$*. If* F *is a destructor, define* $eval_F(\underline{t}) := F(\underline{t})$ *if* $F(\underline{t}) \neq \bot$ *and* $eval_F(\underline{t}) := \bot$ *otherwise.*

Protocols. In CoSP, protocols are represented as infinite trees with the following nodes: *computation nodes* are used for drawing fresh nonces and for applying constructors and destructors; *input* and *output nodes* are used for send and receive operations; *control nodes* are used for allowing the attacker to schedule the protocol. A computation node is annotated with its arguments and has two outgoing edges: a yes-edge, used for the application of constructors, for drawing a nonce, and for the successful application of a destructor, and a no-edge, used if an application of a constructor or destructor F on a term t fails, i.e., if $eval_F(t) = \bot$. Nodes have explicit *references* to other nodes whose terms they use. For example, a computation node that computes $C(t)$ references the node that produced t, e.g., an input node or another computation node.

Definition 4 (CoSP protocol)**.** *A* CoSP protocol Π *for a symbolic model* $\mathbf{M} = (\mathbf{C}, \mathbf{N}, \mathbf{T}, \mathbf{D})$ *is a tree of infinite depth with a distinguished root and labels on both edges and nodes. Each node has a unique identifier* ν *and one of the following types:*[1]

[1] In contrast to the definition in the original CoSP framework [6], we do not consider the type "non-deterministic node". This type has been obsolete for all embeddings of symbolic calculi that have been established in the original framework so far as well.

- Computation nodes *are annotated with a constructor, nonce or destructor* $F/n \in \mathbf{C} \cup \mathbf{N_P} \cup \mathbf{D}$ *together with the identifiers of* n *(not necessarily distinct) nodes; we call these annotations* references*, and we call the referenced nodes* arguments*. Each computation node* ν *has exactly two successors;* $yes(\nu)$ *and* $no(\nu)$*, the corresponding edges are labeled with* `yes` *and* `no`*, respectively.*
- Input nodes *have no annotations. They have exactly one successor,* $succ(\nu)$*.*
- Output nodes *have a reference to exactly one node in their annotations. They have exactly one successor,* $succ(\nu)$*.*
- Control nodes *are annotated with a bitstring* l*. They have at least one and up to countably many successors; the corresponding edges are labeled with distinct bitstrings* l'*. We call* l *the* out-metadata *and* l' *the* in-metadata *of this node.*

We assume that the annotations are part of the node identifier. A node ν *can only reference other nodes* ν' *on the path from the root to* ν*; in this case* ν' *must be a computation node or input node. If* ν' *is a computation node, the path from* ν' *to* ν *has additionally to go through the outgoing edge of* ν' *with label* `yes`*.*

Symbolic operations. As a next step, we model the capabilities of the symbolic attacker. We have to capture which protocol messages the attacker observes, in particular in which order an attacker observe these messages. Moreover, we have to capture which tests an attacker can perform in order to judge whether two protocols are distinguishable. These tests, called *symbolic operations*, capture the sequence of operations a symbolic attacker applies, including the used protocol messages. This notion is sometimes referred to as a *recipe*.

Definition 5 (Symbolic Operation)**.** *Let* $\mathbf{M} = (\mathbf{C}, \mathbf{N}, \mathbf{T}, \mathbf{D})$ *be a symbolic model. A* symbolic operation O/n (of arity n) on $\mathbf{M}$ *is a finite tree whose nodes are labeled with constructors from* $\mathbf{C}$*, destructors from* $\mathbf{D}$*, nonces from* $\mathbf{N}$*, and formal parameters* x_i *with* $i \in \{1, \ldots, n\}$*. For constructors and destructors, the children of a node represent its arguments (if any). Formal parameters* x_i *and nonces do not have children. The set of symbolic operations of arity* n *for a model* $\mathbf{M}$ *is denoted as* $\mathrm{SO}(\mathbf{M})_n$*. The set* $\mathrm{SO}(\mathbf{M}) := \bigcup_{n \in \mathbb{N}} \mathrm{SO}(\mathbf{M})_n$ *is the set of all symbolic operations for* $\mathbf{M}$*. We extend the evaluation function to symbolic operations. Given a list of terms* $\underline{t} \in \mathbf{T}^n$*, the evaluation function* $eval_O : \mathbf{T}^n \to \mathbf{T}$ *recursively evaluates the tree* O *starting at the root as follows: The formal parameter* x_i *evaluates to* t_i*. A node with* $F \in \mathbf{C} \cup \mathbf{N}_E \cup \mathbf{D}$ *evaluates according to* $eval_F$*, where* $\mathbf{N}_E \subseteq \mathbf{N}$ *are attacker nonces. If there is a node that evaluates to* $\bot$*, the whole tree evaluates to* $\bot$*.*

Note that the identity function is included. It is the tree that contains only x_1 as node.

Symbolic execution. A symbolic execution is a path through a protocol tree defined as defined below. It induces a *symbolic view*, which contains the communication with the attacker. We, moreover, define an *attacker strategy* as the sequence of symbolic operations that the attacker performs in the symbolic execution. It induces a *view*, which contains the communication with the attacker. Together with the symbolic execution, we define an *attacker strategy* as the sequence of symbolic operations that the attacker performs.

Definition 6 (Symbolic executions)**.** *Let a symbolic model* $\mathbf{M} = (\mathbf{C}, \mathbf{N}, \mathbf{T}, \mathbf{D})$ *and a CoSP protocol* Π *for* $\mathbf{M}$ *be given. Let* $Views = (Event_{in} \cup Event_{out} \cup Event_{ctl})^*$*, with* $Event_{in} := \{\mathtt{in}\} \times \mathbf{T} \times \mathrm{SO}(\mathbf{M})$*,* $Event_{out} := \{\mathtt{out}\} \times \mathbf{T}^*$*,* $Event_{ctl} := \{\mathtt{control}\} \times \{0,1\}^* \times \{0,1\}^*$*. We define*

$$\rightsquigarrow_{CoSPs} \subseteq Views^* \times Nodes \times (Nodes \to \mathbf{T}) \to Views^* \times Nodes \times (Nodes \to \mathbf{T})$$

as the smallest relation s.t. the rules from Figure 2 hold, where $\tilde{\underline{t}} := f(\tilde{\nu}_1), \ldots, f(\tilde{\nu}_{|\underline{\tilde{\nu}}|})$*, for the nodes* $\tilde{\nu}_1, \ldots, \tilde{\nu}_{|\underline{\tilde{\nu}}|}$ *referenced by* ν *(for computation nodes). The set of* symbolic executions of Π *is defined*

$(V,\nu,f) \rightsquigarrow_{CoSPs} (V, yes(\nu), f[\nu \mapsto m])$	ν computation n. with $F \in \mathbf{C} \cup \mathbf{D} \cup \mathbf{N}$, $m := eval_F(\tilde{t}) \neq \bot$
$(V,\nu,f) \rightsquigarrow_{CoSPs} (V, no(\nu), f)$	ν computation node with $F \in \mathbf{C} \cup \mathbf{D} \cup \mathbf{N}$, $eval_F(\tilde{t}) = \bot$
$(V,\nu,f) \rightsquigarrow_{CoSPs} (V \cdot (\mathtt{in}, (t, O)), succ(\nu), f[\nu \mapsto t])$	ν input node, $t \in \mathbf{T}$, $O \in \mathrm{SO}(\mathbf{M})$, $eval_O(Out(V)) = t$
$(V,\nu,f) \rightsquigarrow_{CoSPs} (V \cdot (\mathtt{out}, \tilde{t}_1), succ(\nu), f)$	ν output node
$(V,\nu,f) \rightsquigarrow_{CoSPs} (V \cdot (\mathtt{control}, (l, l')), \nu', f)$	ν control n., out-metadata l, successor ν' has in-metadata l'

Figure 2: Rules for defining the smallest relation for the symbolic execution

as

$$\begin{aligned}\mathrm{SExec}(\Pi) :=& \{((V_0,\nu_0,f_0),\ldots,(V_n,\nu_n,f_n)) \mid n \in \mathbb{N} \\ & \wedge \forall i.(V_i,\nu_i,f_i) \rightsquigarrow_{CoSPs} (V_{i+1},\nu_{i+1},f_{i+1})\},\end{aligned}$$

where $(V_0,\nu_0,f_0) = (\epsilon, root(\Pi), \emptyset)$. V_i *is called a* symbolic view *for step* i. *The set of symbolic views of* Π *is defined as*

$$\mathrm{SViews}(\Pi) := \{V_n | (V_0,\nu_0,f_0),\cdots,(V_n,\nu_n,f_n) \in \mathrm{SExec}(\Pi)\}$$

Given a view V, $Out(V)$ *denotes the list of terms* t *contained in* $(\mathtt{out}, t) \in V$. $Out\text{-}Meta(V)$ *denotes the list of terms* l *contained in elements of the form* $(\mathtt{control}, (l, l'))$ *in the view* V. $In(V)$, *called the* attacker strategy, *denotes the list of terms that contains only entries of* V *of the form* $(\mathtt{in}, (t, O))$ *or* $(\mathtt{control}, (l, l'))$, *where for* $(\mathtt{in}, (t, O))$ *the entry* $(\mathtt{in}, O)$ *is stored and for* $(\mathtt{control}, (l, l'))$ *the entry* $(\mathtt{control}, l')$ *is stored.* $[In(V)]_{\mathrm{SViews}(\Pi)}$ *denotes the equivalence class of all views* $U \in \mathrm{SViews}(\Pi)$ *with* $In(U) = In(V)$.

Symbolic knowledge. The *symbolic knowledge* of the attacker comprises the results of all symbolic tests the attacker can perform on the messages output by the protocol. The definition captures that the attacker knows exactly which symbolic operation leads to which result.

Definition 7 (Symbolic knowledge)**.** *Let* $\mathbf{M}$ *be a symbolic model. Given a view* V *with* $|Out(V)| = n$, *a* symbolic knowledge function $f_V : \mathrm{SO}(\mathbf{M})_n \to \{\top, \bot\}$ *is a partial function from symbolic operations (see Definition 5) of arity* n *to* $\{\top, \bot\}$. *The* full symbolic knowledge function *is a total symbolic knowledge function* $K_V(O)$ *defined by*

$$K_V(O) := \begin{cases} \bot & \text{if } eval_O(Out(V)) = \bot \\ \top & \text{otherwise.} \end{cases}$$

3.2 Symbolic Equivalence

Two views are *equivalent*, denoted as $V \sim V'$, if they (i) have the same structure (i.e., the same order of $\mathtt{out}, \mathtt{in}, \mathtt{control}$ entries), (ii) have the same out-metadata (i.e., $V_{Out\text{-}Meta} = V'_{Out\text{-}Meta}$), and (iii) lead to the same knowledge (i.e., $K_V = K_{V'}$). Finally, two CoSP protocols are *symbolically equivalent* ($\approx_s^{\mathrm{CoSP}}$) if its two variants lead to equivalent views when run with the same attacker.

Definition 8 (Equivalent views)**.** *Let two views* V, V' *of the same length be given. We denote their* i*th entry by* V_i *and* V'_i, *respectively.* V *and* V' *are* equivalent *(*$V \sim V'$*), if the following three conditions hold:*

1. *(Same structure)* V_i *is of the form* $(s, \cdot)$ *if and only if* V'_i *is of the form* $(s, \cdot)$ *for some* $s \in \{\mathtt{out}, \mathtt{in}, \mathtt{control}\}$.

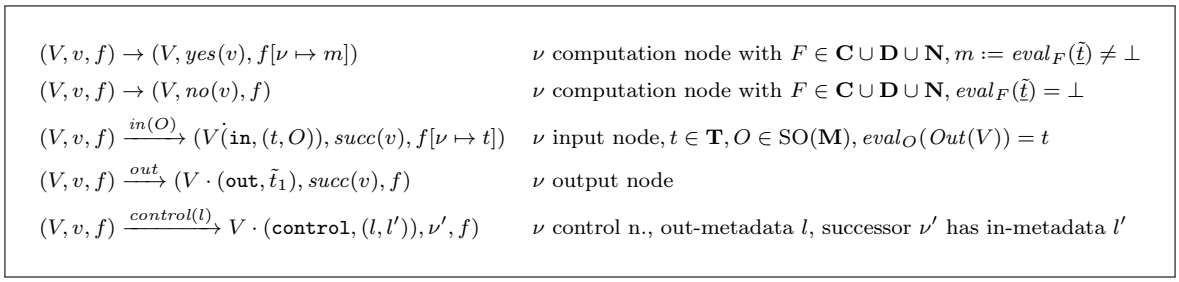
$(V, v, f) \to (V, yes(v), f[\nu \mapsto m])$ ν computation node with $F \in \mathbf{C} \cup \mathbf{D} \cup \mathbf{N}, m := eval_F(\tilde{t}) \neq \bot$
$(V, v, f) \to (V, no(v), f)$ ν computation node with $F \in \mathbf{C} \cup \mathbf{D} \cup \mathbf{N}, eval_F(\tilde{t}) = \bot$
$(V, v, f) \xrightarrow{in(O)} (V \dot{(}\mathtt{in}, (t, O)), succ(v), f[\nu \mapsto t])$ ν input node, $t \in \mathbf{T}, O \in \mathrm{SO}(\mathbf{M}), eval_O(Out(V)) = t$
$(V, v, f) \xrightarrow{out} (V \cdot (\mathtt{out}, \tilde{t}_1), succ(v), f)$ ν output node
$(V, v, f) \xrightarrow{control(l)} V \cdot (\mathtt{control}, (l, l')), \nu', f)$ ν control n., out-metadata l, successor ν' has in-metadata l'

Figure 3: Labelled transition

2. *(Same out-metadata)* $\mathit{Out\text{-}Meta}(V) = \mathit{Out\text{-}Meta}(V')$.

3. *(Same symbolic knowledge)* $K_V = K_{V'}$.

Finally, we define two protocols to be *symbolically equivalent* if the two protocols lead to equivalent views when faced with the same attacker strategy. Thus, a definition of symbolic equivalence should compare the symbolic knowledge of two protocol runs only if the attacker behaves identically in both runs.

Given the static equivalence relation from Definition 8, the notion of symbolic equivalence in CoSP is defined as trace equivalence.

Definition 9 (Symbolic equivalence)**.** *Let* $\mathbf{M}$ *be a symbolic model and* P *be a class of protocols for* $\mathbf{M}$*. Let* $\mathit{traces}_A(\Pi) = \left\{[A]_{\mathrm{SViews}(\Pi)}\right\}$ *denote the list of traces for an attacker strategy* A*. We lift the relation* $\sim$ *from Definition 8 to sets, i.e.,* $A \sim B$ *iff.* $\forall a \in A \exists b \in B. a \sim b$ *and vice versa.*

Two protocols $\Pi_1, \Pi_2 \in \mathrm{P}$ *are* trace equivalent *(*$\Pi_1 \approx_s^{\mathrm{CoSP}} \Pi_2$*), if, for all attacker strategies* A*,*

$$\mathit{traces}_A(\Pi_1) \sim \mathit{traces}_A(\Pi_2).$$

Definition 10 ((Weak) bisimilarity)**.** *A* (weak) bisimulation relation *is a symmetric relation*

$$\mathcal{R} : \mathit{Views}^* \times \mathit{Nodes} \times (\mathit{Nodes} \to \mathbf{T}) \times \mathit{Views}^* \times \mathit{Nodes} \times (\mathit{Nodes} \to \mathbf{T})$$

such that, if $A \mathcal{R} B$*, then*

- $K_{V_A} = K_{V_B}$ *(where* V_A*,*V_B *are the symbolic views in* A *and* B*).*
- *if* $A \to A'$ *then there is* B' *s.t.* $B \to^* B'$ *and* $A \mathcal{R} B$
- *if* $A \xrightarrow{\alpha} A'$ *then there is* B' *such that* $B \to^* \xrightarrow{\alpha} \to^* B'$ *and* $A \mathcal{R} B$*.*

If there exist a weak bisimulation relation between A *and* B*, they are* bisimilar*.*

Without non-deterministic nodes, more generally without internal non-determinism, the notion of trace equivalence (i.e., symbolic equivalence from Definition 9) is equivalent to weak bisimilarity.

Lemma 1. *For any model* $\mathbf{M}$*, any two processes* Π_1 *and* Π_2 *are trace equivalent if they are bisimilar.*

Proof. We introduce the following function mapping a single stap of a trace, i.e., $e \in \mathit{Views}$ to labels according to Definition 3. Let

$$\beta(v) = \begin{cases} \mathtt{in}(O) & \text{if } v = (\mathtt{in}, (t, O)) \in \mathit{Event}_{in} \\ \mathtt{out} & \text{if } v = (\mathtt{out}, \tilde{t}) \in \mathit{Event}_{out} \\ \mathtt{control}(l) & \text{if } v = (\mathtt{control}(l, l')) \in \mathit{Event}_{ctl}. \end{cases}$$

$(\nu, f, n) \rightsquigarrow_{CoSPc} (yes(\nu), f[\nu \mapsto n(N)], n)$	ν computation node with $N \in \mathbf{N}$ and $n(N) \neq \bot$
$(\nu, f, n) \rightsquigarrow_{CoSPc} (yes(\nu), f[\nu \mapsto m], n[N \mapsto m])$	ν computation node with $N \in \mathbf{N}$ and $n(N) = \bot$ and m drawn according to $\mathrm{Impl}_N(\eta)$.
$(\nu, f, n) \rightsquigarrow_{CoSPc} (yes(\nu), f[\nu \mapsto m], n)$	ν computation node with $F \in \mathbf{C} \cup \mathbf{D} \cup \mathbf{N}$, and $A_F(\eta, \underline{\tilde{t}}) = m \neq \bot$
$(\nu, f, n) \rightsquigarrow_{CoSPc} (no(\nu), f, n)$	ν computation node with $F \in \mathbf{C} \cup \mathbf{D} \cup \mathbf{N}$, and $A_F(\eta, \underline{\tilde{t}}) = \bot$
$(\nu, f, n) \rightsquigarrow_{CoSPc} (succ(\nu), f[\nu \mapsto m], n)$	ν input node, a request bitstring and receive m
$(\nu, f, n) \rightsquigarrow_{CoSPc} (succ(\nu), f, n)$	ν output note, send $\tilde{m}_1$ to the communication partner
$(\nu, f, n) \rightsquigarrow_{CoSPc} (\nu', f, n)$	ν control node with out-metadata l, send l to communication partner and receive in-metadata l'. If there is an edge with label l', let ν' be the successor of ν along the edge labeled l'; otherwise let ν' be the successor with the lexicographically smallest edge label.

Figure 4: Rules defining the smallest relation for the computational execution, where $yes(\nu)$ and $no(\nu)$ are the yes- and no-successor of ν (for computation nodes), *succ* is the successor node (for input or output nodes), and $\underline{\tilde{t}} := \tilde{t}_j := f(\tilde{\nu}_j)$ with $\underline{\tilde{\nu}}$ are the nodes referenced by ν (for computation nodes).

Without loss of generality, we assume that each CoSP tree is transformed into a CoSP tree where each node uniquely describes its subtree completely. Fix an arbitrary trace $t_1 = (v_1, \ldots, v_n) \in traces_A(\Pi_1)$. By Definition 6 and Figure 3, there is a sequence of transitions and triples (V_i, ν_i, f_i) such that:

$$(V_0, \nu_0, f_0) \rightarrow^* \xrightarrow{\beta(v_1)} \rightarrow^* \cdots \rightarrow^* \xrightarrow{\beta(v_n)} \rightarrow^* (V_n, \nu_n, f_n),$$

and $v_0 = \Pi_1$. As Π_1 and Π_2 are bisimilar, there exists a sequence of transitions and triples (V_i', ν_i', f_i') such that $v_0' = \Pi_2$, and

$$(V_0', \nu_0', f_0') \rightarrow^* \xrightarrow{\beta(v_1)} \rightarrow^* \cdots \rightarrow^* \xrightarrow{\beta(v_n)} \rightarrow^* (V_n', \nu_n', f_n'),$$

as well as $K_{V_n} = K_{V_{n'}}$. The latter point implies that the third condition of Definition 8 holds for $(V_0, \ldots, V_n)$ and $(V_0', \ldots, V_n')$. From the relation between the view and the labels in each transition described in Definiton 3, it follows that the first and second condition holds. Hence for each $(V_0, \ldots, V_n) \in \mathrm{SViews}(\Pi_1)$ there exists an equivalent view $(V_0', \ldots, V_n') \in \mathrm{SViews}(\Pi_2)$. The same argument can be made for Π_2 in place of Π_1 and vice versa, thus Π_1 and Π_2 are trace equivalent.

□

3.3 Computational execution

In the computational setting, symbolic constructors and destructors are realized with cryptographic algorithms. A *computational implementation* is a family $\underline{\mathrm{Impl}} = (A_x)_{x \in \mathbf{C} \cup \mathbf{D} \cup \mathbf{N_P}}$ of deterministic polynomial-time algorithms Impl_F for each constructor or destructor $F \in \mathbf{C} \cup \mathbf{D}$ as well as a probabilistic polynomial-time (ppt) algorithm A_N for drawing protocol nonces $N \in \mathbf{N}$. The first argument of Impl_F and A_N represents the security parameter.

Definition 11 (Computational implementation)**.** *Let* $\mathbf{M} = (\mathbf{C}, \mathbf{N}, \mathbf{T}, \mathbf{D})$ *be a symbolic model. A* computational implementation of $\mathbf{M}$ *is a family of functions* $\mathrm{Impl} = (A_x)_{x \in \mathbf{C} \cup \mathbf{D} \cup \mathbf{N}}$ *such that* A_F *for* $F/n \in \mathbf{C} \cup \mathbf{D}$ *is a partial deterministic function* $\mathbb{N} \times (\{0,1\}^*)^n \to \{0,1\}^*$*, and* A_N *for* $N \in \mathbf{N}$ *is a total probabilistic function with domain* $\mathbb{N}$ *and range* $\{0,1\}^*$*. All functions* A_F *have to be computable in deterministic polynomial time, and all* A_N *have to be computable in probabilistic polynomial time (ppt).*

We extend the computational implementation to symbolic operations. The computational implementation of a symbolic operation $O \in \mathrm{SO}$ *is defined* $A_O : (\{0,1\}^*)^n \to \{0,1\}^*$*, recursively evaluating*

the tree O starting at the root as follows: The formal parameter x_i computed to ith input value. A node with $F \in \mathbf{C} \cup \mathbf{N}_E \cup \mathbf{D}$ is computed according to A_F. If the function is undefined for any node, then the A_O is undefined on these inputs.

Computational execution. The *computational execution* of a protocol is a randomized interactive machine, called the *computational challenger*, that traverses the protocol tree and interacts with a ppt attacker $\mathcal{A}$: at a computation node the corresponding algorithm is run and depending on whether the algorithm succeeds or outputs $\perp$, either the yes-branch or the no-branch is taken; at an output node, the message is sent to the attacker, and at an input node a message is received by the attacker; at a control node the attacker sends a command that specifies which branch to take. The transcript of the execution contains the computational counterparts of a symbolic view.[2]

Definition 12 (Computational execution)**.** *Let* Impl *be a computational implementation of the symbolic model $\mathbf{M} = (\mathbf{C}, \mathbf{N}, \mathbf{T}, \mathbf{D})$ and Π be a CoSP protocol. For a security parameter η, the* computational challenger $\mathrm{Exec}_{\mathbf{M},\mathrm{Impl},\Pi}(\eta)$ *is an interactive Turing machine defined by repeatedly applying the rules in Figure 4, starting from the initial state (ν, f, n), where ν is the root of Π, and f and n are undefined partial functions from node identifiers to bitstrings and from $\mathbf{N}$ to bitstrings, respectively.*

We call the interaction between $\mathrm{Exec}_{\mathbf{M},\mathrm{Impl},\Pi}(\eta)$ and an interactive ppt Turing machine $\mathcal{A}(\eta)$ the computational execution*, and denote it as $\langle \mathrm{Exec}_{\mathbf{M},\mathrm{Impl},\Pi}(\eta) \| \mathcal{A}(\eta) \rangle$ using standard notation for interacting machines. The execution stops whenever one of the two machines stops, and the output of $\langle \mathrm{Exec}_{\mathbf{M},\mathrm{Impl},\Pi}(\eta) \| \mathcal{A}(\eta) \rangle$ is the output of $\mathcal{A}(\eta)$.*

By Definition 19, there is an efficiently computable injection from $\mathcal{V} \cup \mathcal{H}$ into $\{0,1\}^*$, hence we abuse notation by supplying values from domain in lieu of bitstrings, and using the inverse of this injection to convert bitstrings back into values from $\mathcal{V} \cup \mathcal{H}$.

The CoSP execution is used in cryptographic reductions; however, CoSP protocols are by definition infinite. Hence, as is standard in cryptography, we have to ensure the execution of a CoSP protocols is computable in polynomial-time. To this end, we require that it is possible to incrementally compute the node information, i.e., the node identifier, for each path through the CoSP protocol in polynomially many steps (in the length of the path from the root node to the current node). We call such (bi-)protocols efficient.

Definition 13 (Efficient Protocol)**.** *We call a CoSP protocol* efficient *if:*

- *There is a polynomial p such that for any node N, the length of the identifier of N is bounded by $p(m)$ where m is the length (including the total length of the edge-labels) of the path from the root to N.*
- *There is a deterministic polynomial-time algorithm that, given the identifiers of all nodes and the edge labels on the path to a node N, computes the identifier of N.*
- *There is a deterministic polynomial-time algorithm that, given the identifier of a control node N, the identifiers of all nodes and all edge labels on the path to N, computes the lexicographically smallest label of an edge (i.e., the in-metadata) of all edges that lead from N to one of its successors.*

Computational indistinguishability. For computational soundness results of equivalence properties in CoSP, the notion of computational indistinguishability that is prevalently used in cryptography has to be refined to the notion of *termination-insensitive computational indistinguishability* (tic-indistinguishability, in short) [12]. In comparison to the standard notion of indistinguishability, tic-indistinguishability does not require the interactive machines to be polynomial-time, but it instead

[2] We stress that we not have to incorporate an explicit run-time polynomial to ensure overall termination in polynomial time, since computational soundness in CoSP is based on tic-indistinguishability, which only requires indistinguishability for polynomially-long prefixes, see Section 3.4.

only considers decisions that were made for polynomially-bounded prefixes of the interaction. This excludes pathological cases in which programs of super-polynomial running time only differ in the response time (which an adversary might not be able to react to in polynomial time) and in which cryptographic operations might be executed on inputs of super-polynomial length (for which common cryptographic definitions do not give any guarantees).

Definition 14 (Tic-indistinguishability)**.** *Given two machines M, M' and a polynomial p, we write $\Pr[\langle M \| M' \rangle \downarrow_{p(\eta)} x]$ for the probability that the interaction between M and M' terminates within $p(\eta)$ steps and M' outputs x.*

We call two machines A and B termination-insensitively computationally indistinguishable for a machine $\mathcal{A}$ *($A \approx^{\mathcal{A}}_{tic} B$) if for all polynomials p, there is a negligible function μ such that for all $z, a, b \in \{0,1\}$ with $a \neq b$,*

$$\Pr[\langle A(\eta) \| \mathcal{A}(\eta, z) \rangle \downarrow_{p(\eta)} a] + \Pr[\langle B(\eta) \| \mathcal{A}(\eta, z) \rangle \downarrow_{p(\eta)} b] \leq 1 + \mu(\eta).$$

Here, z represents an auxiliary string. Additionally, we call A and B termination-insensitively computationally indistinguishable *($A \approx_{tic} B$) if we have $A \approx^{\mathcal{A}}_{tic} B$ for all polynomial-time machines $\mathcal{A}$.*

For two CoSP protocols Π_1, Π_2, a symbolic Model $\mathbf{M}$, and implementation $\mathbf{Impl}$, if $\mathbf{M}$ and $\mathbf{Impl}$ are clear from the context, we use for $\mathrm{Exec}_{\mathbf{M},\mathrm{Impl},\Pi_1}(\eta) \approx_{tic} \mathrm{Exec}_{\mathbf{M},\mathrm{Impl},\Pi_2}(\eta)$ (for sufficiently large $k \in \mathbb{N}$) the notation $\Pi_1 \approx^{\mathrm{CoSP}}_{c} \Pi_2$.

3.4 Computational Soundness

The previous notions culminate in the definition of computational soundness for equivalence properties. It states that the symbolic equivalence of efficient protocols implies their computational indistinguishability.

Definition 15 (Computational soundness)**.** *Let a symbolic model $\mathbf{M}$ and a class* P *of efficient protocols be given. A computational implementation $\mathbf{Impl}$ of $\mathbf{M}$ is* computationally sound *for $\mathbf{M}$ if for every pair $\Pi_1, \Pi_2 \in$ P symbolic equivalence $\Pi_1 \approx^{\mathrm{CoSP}}_{s} \Pi_2$ implies $\Pi_1 \approx^{\mathrm{CoSP}}_{c} \Pi_2$, i.e., tic-indistinguishability (see Definition 14) of their computational execution $\mathrm{Exec}_{\mathbf{M},\mathrm{Impl},\Pi_1}(\eta) \approx_{tic} \mathrm{Exec}_{\mathbf{M},\mathrm{Impl},\Pi_2}(\eta)$.*

4 Dalvik Bytecode (Review)

The Abstract Dalvik Language (ADL) [4] constitutes the currently most detailed and comprehensive operational semantics for Dalvik in the literature, even though it does not encompass concurrency and exceptions (as e.g. in [3]). We refer to Section 8 for further discussion on different Dalvik semantics. We only provide a compact review of ADL and refer to [4] for more details.

4.1 Syntax of ADL

ADL uses as syntactical domains seven underspecified sets and several sets of unary, binary, and relation operations.

Definition 16 (Syntactical domains of ADL)**.** *The following sets constitute the* syntactical domains of ADL.

$$
\begin{array}{rll}
\textsc{rConst}\colon & s \to s\{pp+1, r[v_a \mapsto n]\} & \text{for } m[pp] = \mathsf{const}\ v_a, n \\
\textsc{rMove}\colon & s \to s\{pp+1, r[v_a \mapsto r(v_b)]\} & \text{for } m[pp] = \mathsf{move}\ v_a, v_b \\
\textsc{rUnop}\colon & s \to s\{pp+1, r[v_a \mapsto u]\} & \text{for } m[pp] = \mathsf{unop}\ v_a, v_b, uop \text{ and } u = \underline{uop}(r(v_b)) \\
\textsc{rBinop}\colon & s \to s\{pp+1, r[v_a \mapsto u]\} & \text{for } m[pp] = \mathsf{binop}\ v_a, v_b, v_c, bop \text{ and } u = r(v_b)\ \underline{bop}\ r(v_c) \\
\textsc{IfTestT}\colon & s \to s\{pp+n\} & \text{for } m[pp] = \mathsf{if\text{-}test}\ v_a, v_b, n, rop \text{ and } r(v_a)\ \underline{rop}\ r(v_b) \\
\textsc{IfTestF}\colon & s \to s\{pp+1\} & \text{for } m[pp] = \mathsf{if\text{-}test}\ v_a, v_b, n, rop \text{ and } \neg(r(v_a)\ \underline{rop}\ r(v_b)) \\
\textsc{rMoveR}\colon & s \to s\{pp+1, r[v_a \mapsto r(res_{lo})]\} & \text{for } m[pp] = \mathsf{move\text{-}result}\ v_a \\
\textsc{rISt}\colon & s \to s\left\{\begin{array}{l} m' \cdot m \cdot ml, 0 \cdot pp \cdot ppl, \\ defReg([r(v_a), \ldots, r(v_e)]) \cdot r \cdot rl \end{array}\right\} & \begin{array}{l}\text{for } m[pp] = \mathsf{invoke\text{-}static}\ v_a, \ldots, v_e, mid \\ \quad m' = lookup\text{-}static_{\Pi}(mid)\end{array} \\
\textsc{rIDR}\colon & s \to s\left\{\begin{array}{l} m' \cdot m \cdot ml, 0 \cdot pp \cdot ppl, \\ defReg([r(v_k), \ldots, r(v_{k+n-1})]) \cdot r \cdot rl \end{array}\right\} & \begin{array}{l}\text{for } m[pp] = \mathsf{invoke\text{-}direct\text{-}range}\ v_k, n, mid \\ \quad m' = lookup\text{-}direct_{\Pi}(mid, h(r(v_k)).\mathsf{class})\end{array} \\
\textsc{rReturnV}\colon & \langle m \cdot ml, h, pp \cdot pp' \cdot ppl, r \cdot r' \cdot rl, as\rangle \to & \text{for } m[pp] = \mathsf{return\text{-}void} \text{ and } ml' \neq () \\
& \langle ml, h, (pp'+1) \cdot ppl, & \\
& r'[res_{lo} \mapsto lo(\text{void}), res_{up} \mapsto up(\text{void})] \cdot rl, as\rangle & \\
\textsc{rReturn}\colon & \langle m \cdot ml, h, pp \cdot pp' \cdot ppl, r \cdot r' \cdot rl, as\rangle \to & \text{for } m[pp] = \mathsf{return}\ v_a \text{ and } ml' \neq () \\
& \langle ml, h, (pp'+1) \cdot ppl, r'[r''] \cdot rl, as\rangle & \quad r'' = res_{lo} \mapsto lo(r(v_a)), res_{up} \mapsto up(r(v_a)) \\
\textsc{rReturnVF}\colon & \langle (m), h, (pp), (r), as\rangle \to \langle \text{void}, h, as\rangle & \text{for } m[pp] = \mathsf{return\text{-}void} \\
\textsc{rReturnF}\colon & \langle (m), h, (pp), (r), as\rangle \to \langle r(v_a), h, as\rangle & \text{for } m[pp] = \mathsf{return}\ v_a
\end{array}
$$

for $v_a, v_b, v_c, v_d, v_e \in \mathcal{X}$, $n \in \mathbb{N}$, $mid \in \mathcal{MID}$, $uop \in \mathcal{UNOP}$, $bop \in \mathcal{BINOP}$, $rop \in \mathcal{RELOP}$ and $u \in \mathcal{N}$.

Figure 5: Selection of inference rules that define $\to$ with $uop \in \mathcal{UNOP}, bop \in \mathcal{BINOP}, rop \in \mathcal{RELOP}$, with some $n \in \mathbb{N}$ counter for method calls, some ADL program Π, and some method m. We write $v_a, \ldots, v_e$ for v_a, v_b, v_c, v_d, v_e, and if $r = defReg(u_1, \ldots, u_l)$, then $r(v_i) = u_i$ for $i \in \{0, \ldots, l\}$ and void otherwise.

$$
\begin{array}{rl}
\mathcal{CONV} & \textit{: set of symbols for type-casting operations} \\
\mathcal{UNOP} & := \{-, \neg\} \cup \mathcal{CONV} \textit{ unary operations} \\
\mathcal{BINOP} & := \{+, -, *, /, \%, \wedge, \vee, \oplus, <<, >>, >>>\} \\
 & \textit{binary operations} \\
\mathcal{RELOP} & := \{=, \neq, <, >, \leq, \geq\} \textit{ relations} \\
\mathcal{X} & := \{v_i \mid i \in \mathbb{N}\} \textit{ register names}
\end{array}
\qquad
\begin{array}{rl}
\multicolumn{2}{l}{\textit{underspecified sets}} \\
\mathcal{CID} & \textit{: set of class names} \\
\mathcal{FID} & \textit{: set of field names} \\
\mathcal{MID} & \textit{: set of method names} \\
\mathcal{S} & \textit{: set of string symbols} \\
\mathcal{N} & \textit{: set of numerical symbols} \\
 & \textit{(e.g., integers and floating point numbers)} \\
\mathcal{L}_c & \textit{: set of constant memory locations} \\
\mathcal{F} & \textit{: set of fields}
\end{array}
$$

We require that $\mathcal{MID}$,$\mathcal{CID}$, and $\mathcal{FID}$ are mutually disjoint. A given program typically only uses a subset of the underspecified sets. In these cases, we use the notation K_y to denote that K_y is some subset of K ($K_y \subseteq K$) specific to y, for all $K \in \{\mathcal{CID}, \mathcal{FID}, \mathcal{MID}, \mathcal{S}, \mathcal{N}, \mathcal{F}\}$.

Many ADL instructions have a -wide variant for 2-register-values and a -range variant where instead of 5 registers only the starting register and a range n is given and the arguments are read of the n subsequent registers, beginning from the starting register. Figure 6 depicts the set of ADL instructions but omits all -wide and -range variants for the sake of readability, as they are treated analogously. The complete set of instructions can be found in [4].

With the syntactical domains (Definition 16) and the set of instructions (Figure 6) at hand, we now define the syntax of ADL methods and of ADL programs.

Definition 17 (ADL methods)**.** *The set $\mathcal{M}$ of* ADL methods *is defined by $\mathcal{M} := \mathcal{INSTR}^* \setminus []$, where $\mathcal{INSTR}$ is the set of instructions (Figure 6).*

ADL uses five partial *lookup functions*. These lookup functions refer to the method names that return the field (*lookup-field*$_\Pi$) and the instruction set with respect to the method's name: for static

Arithmetic instructions	Control flow instructions	Method-related instructions
move v_a, v_b	nop	invoke-virtual $v_a, v_b, v_c, v_d, v_e, n, mid$
const v_a, n	goto n	invoke-super $v_a, v_b, v_c, v_d, v_e, n, mid$
cmp v_a, v_b, v_c	if-test v_a, v_b, n, rop	invoke-direct $v_a, v_b, v_c, v_d, v_e, n, mid$
unop v_a, v_b, uop	if-testz v_a, n, rop	invoke-interface $v_a, \ldots, v_e, n, mid$
binop v_a, v_b, v_c, bop	Object-related instructions	invoke-static $v_a, v_b, v_c, v_d, v_e, n, mid$
	instance-of v_a, v_b, cl	move-result v_a
Array-related instructions	new-instance v_a, cl	return-void
array-length v_a, v_b	const-string v_a, s	return v_a
new-array v_a, v_b	const-class v_a, cl	
filled-array $v_a, v_b, v_c, v_d, v_e, n$	iget v_a, v_b, fid	
fill-array-data $v_a, u_0, \ldots, u_n$	iput v_a, v_b, fid	
aget v_a, v_b, v_c	sget v_a, fid	
aput v_a, v_b, v_c	sput v_a, fid	

for $v_a, v_b, v_c, v_d, v_e \in \mathcal{X}$, $n \in \mathbb{N}$, $mid \in \mathcal{MID}$, $fid \in \mathcal{FID}$, $cl \in \mathcal{CID}$, $uop \in \mathcal{UNOP}$, $bop \in \mathcal{BINOP}$, $rop \in \mathcal{RELOP}$, and $u_i \in \mathcal{N} \cup \mathcal{L}_c$ (for $i \in \mathcal{N}_0$).

Figure 6: The set $\mathcal{INSTR}$ of ADL Instructions (excluding all -wide and -range variants of commands)

methods (*lookup-static*$_\Pi$), normal methods (*lookup-direct*$_\Pi$), inherited methods (*lookup-super*), and virtual methods (*lookup-virtual*$_\Pi$).

Definition 18 (ADL programs)**.** *An* ADL program Π *is a tuple*

$$\begin{aligned}\Pi := \ & (\mathcal{CID}_\Pi, \mathcal{FID}_\Pi, \mathcal{MID}_\Pi, \mathcal{M}_\Pi, \mathcal{F}_\Pi, \\ & \textit{lookup-field}_\Pi, \textit{lookup-static}_\Pi, \textit{lookup-direct}_\Pi, \textit{lookup-super}_\Pi, \textit{lookup-virtual}_\Pi), \textit{where}\end{aligned}$$

$$\mathcal{CID}_\Pi \subseteq_{fin} \mathcal{CID}, \mathcal{FID}_\Pi \subseteq_{fin} \mathcal{FID}, \mathcal{MID}_\Pi \subseteq_{fin} \mathcal{MID}, \mathcal{M}_\Pi \subseteq_{fin} \mathcal{M}, \mathcal{F}_\Pi \subseteq_{fin} \mathcal{F},$$

$$\begin{aligned}\textit{lookup-field}_\Pi &: \mathcal{FID}_\Pi \to \mathcal{F}_\Pi, \\ \textit{lookup-static}_\Pi &: \mathcal{MID}_\Pi \rightharpoonup \mathcal{M}_\Pi, \\ \textit{lookup-direct}_\Pi &: \mathcal{MID}_\Pi \times \mathcal{CID}_\Pi \rightharpoonup \mathcal{MID}_\Pi, \\ \textit{lookup-super}_\Pi &: \mathcal{MID}_\Pi \times \mathcal{CID}_\Pi \rightharpoonup \mathcal{MID}_\Pi, \\ \textit{lookup-virtual}_\Pi &: \mathcal{MID}_\Pi \times \mathcal{CID}_\Pi \rightharpoonup \mathcal{MID}_\Pi\end{aligned}$$

4.2 The Semantics of ADL

In this section, we review ADL's operational semantics as far as needed to understand and state our results; in particular, we only provide a partial list of ADL's instructions. We refer to the original ADL paper for further information [4].

The execution relation defines the operational semantics of ADL. Before we define the execution relation, we introduce the semantical domains of ADL, A *state*, also called an *intermediate state*, in ADL consists of a heap h, a program counter pp, and a set of register values $r_1, \ldots, r_n$, and is extended here by some adversarial state as. Formally defining an ADL state and an ADL program requires several additional semantical domains sets and functions. First, we define the set of *registers*. Let $\mathcal{L} = \mathcal{L}_v \cup \mathcal{L}_c$ denote the set of *memory locations*, where $\mathcal{L}_v$ is an underspecified set of *variable locations*

and $\mathcal{L}_c$ an underspecified set of *constant memory locations*, with $\mathcal{L}_v \cap \mathcal{L}_c = \emptyset$. For the purpose of this work, we assume a total order on $\mathcal{L}$. Let $\mathcal{V} = \mathcal{N} \cup \mathcal{L} \cup \{\text{void}\}$ denote the set of *values*, where $\mathcal{N}$ is an underspecified, but finite, set of *numerical values* and void a distinguished return-value for void-methods. Let $\mathcal{R} = \mathbb{N} \cup \{res_{lo}, res_{up}\} \to \mathcal{V}$ be the set of *registers*, where res_{lo} and res_{up} are reserved registers for return values (split into an upper and a lower part) of method calls. To define the set of heap states, let $\mathcal{AR} = \mathbb{N} \times (\mathbb{N} \rightharpoonup \mathcal{V})$. Then the set $\mathcal{H}$ of *heaps* is defined as $\mathcal{H} := (\mathcal{L} \rightharpoonup (\mathcal{O} \cup \mathcal{AR}))$. Altogether, these notions are referred to as the *semantical domains* of ADL.

Throughout the paper, we distinguish between configurations and states. A configuration describes the state of the program without the attacker. Later, we also exclude the state of the library from the configurations. In ADL, configurations are elements from the set $\mathcal{C}'$. In particular, we distinguish initial states and initial configurations.

Definition 19 (Semantical domains)**.** *The semantical domains of ADL programs are defined by*

$$
\begin{aligned}
\mathcal{L} &= \mathcal{L}_c \cup \mathcal{L}_v && \textit{locations}\\
\mathcal{V} &= \mathcal{N} \cup \mathcal{L} \cup \{\text{void}\} && \textit{values}\\
\mathcal{O} &= \mathcal{CID} \times (\mathcal{F} \rightharpoonup \mathcal{V}) && \textit{objects}\\
\mathcal{AR} &= \mathbb{N} \times (\mathbb{N} \rightharpoonup \mathcal{V}) && \textit{arrays}\\
\mathcal{X}_{res} &= \{res_{lo}, res_{up}\} && \textit{reserved registers}\\
\mathcal{R} &= (\mathcal{X} \cup \mathcal{X}_{res}) \to \mathcal{V} && \textit{register states}\\
\mathcal{H} &= (\mathcal{L} \rightharpoonup (\mathcal{O} \cup \mathcal{AR})) && \textit{heaps}\\
\mathcal{C}' &= \mathcal{M} \times \mathcal{H} \times \mathbb{N} \times \mathcal{R} \times \mathcal{Q} && \textit{intermediate configuration}\\
\mathcal{C} &= \mathcal{M} \times \mathcal{H} \times \mathbb{N} \times \mathcal{R} \times \mathcal{Q} && \textit{intermediate state}\\
\mathcal{C}_{final} &= \mathcal{V} \times \mathcal{H} \times \mathcal{Q} \cup \mathcal{ADVR} && \textit{final state}
\end{aligned}
$$

where $\mathcal{L}_v$ with $\mathcal{L}_v \cap \mathcal{L}_c = \emptyset$ is the set of variable locations, void *is a special value such that* void $\notin (\mathcal{N} \cup \mathcal{L})$*, res_{lo} and res_{up} are special registers such that $res_{lo}, res_{up} \notin \mathcal{X}$, and $\mathcal{Q}$ is the state space of the adversary (see Section 5.2). Moreover, there is an efficiently computable injection from $\mathcal{V} \cup \mathcal{H}$ into bitstrings, which can be efficiently inverted on its range.*

The execution relation $\to$. The operational semantics is defined in terms of an execution relation $\to$ (for an ADL program Π, which we assume fixed in this section). For the sake of illustration, Figure 5 contains a representative selection of the rules defining $\to$. For the full set of rules, we refer to the work of Lortz et al. [4].

We use the following notation to shorten presentation and highlight the modifications applied to the state. For a state $s = \langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \in \mathcal{C}$, we use $s\{pp + 1\}$ to denote $\langle m \cdot ml, h, pp + 1 \cdot ppl, r \cdot rl, as\rangle$. Similarly $s\{r[v \mapsto a]\}$ denotes $\langle m \cdot ml, h, pp \cdot ppl, r[v \mapsto a] \cdot rl, as\rangle$, $s\{h[l \mapsto a]\}$ denotes $\langle m \cdot ml, h[l \mapsto a], pp \cdot ppl, r \cdot rl, as\rangle$, $s\{m'\}$ denotes $\langle m' \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle$, and $s\{as'\}$ denotes $\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as'\rangle$.

The relation defines constant assignment (rConst), copying of register values (rMove), binary operations (rBinop), conditional branching (rIfTestTrue and rIfTestFalse). Moreover, we depict rules for static method evaluation (rISt) and evaluation of final methods (rIStR). Return values are stored in distinct result register $res_{lo}, res_{up} \in \mathcal{X}_{res}$ [3](see rReturn).

We slightly diverge from the characterization of method calls in [4] to capture the total number of computation steps in a run: each transition corresponds to one computation step, and each rule in Figure 5 is annotated accordingly ($\to$ instead of $\to_\Pi$).

[3]Although ADL includes two distinct result registers res_{lo} and res_{up} exist, the upper register res_{up} is only used for the move-result-wide instruction, where the return value is copied in two parts.

5 Security framework

We extend ADL with prbabilistic choices and with a probabilistic polynomial-time attacker that is invoked whenever specific functions are invoked. The modifications to the ADL-semantics are depicted in Figure 7.

5.1 Execution and communication model

ADL as defined in [4] does not support probabilistic choice and hence no generation of cryptographic keys and no executions of cryptographic functions. We thus first extended ADL with a rule that uniformly samples a register value from the set of numerical values $\mathcal{N}$ (see the PROB-rule in Figure 7).

To simplify presentation, we interpret the ADL's semantics as a (generative) probabilistic transition system, assuming probability 1 and number of computation steps 1 for all transitions in Figure 5 and Appendix A. The symbolic variant of ADL, presented in the next section, will simplify the adversary by means of a deduction relation, i.e., non-deterministic choice over all message deducible by the adversary. To be able to capture the non-determinism in the symbolic variant, as well as the probabilism in the computational variant, we chose a model of a probabilistic transition system similar to the model introduced by Vardi under the name *concurrent Markov chains* [13], but recast in terms of a transition labelled system (as opposed to state labelled), and with the restriction that probabilistic choices are always unlabelled, which simplifies the definition of parallel composition, which we will later use to decouple, and substitute attacker as well as library. Hence, the definition of the distribution of traces below applies to any probabilistic transition system, including the split-state compositions from Section 7.1. We furthermore annotate both probabilistic and non-deterministic steps with the number of computation steps in $\mathbb{N}$, to be able to argue about the runtime of a system.

Definition 20 (Probabilistic transition system)**.** *A* probabilistic transition system *is a quadruple* $(\mathrm{S}, s_0, \mathrm{A}, \delta)$ *consisting of*

- *a set of states* S*,*
- *an initial state* $s_0 \in \mathrm{S}$*,*
- *a set of actions* A*, and*
- *a transition function* $\delta \colon S \to \mathcal{D}(S \times \mathbb{N}) \uplus \mathcal{P}(A \times S \times \mathbb{N})$*.*

Given $s \in \mathrm{S}$ *and* $\delta(s) = \mu$*, we write* $s \xrightarrow{[p]}_n s'$ *if* $\mu(s', n) = p$*. If* $(a, s', n) \in \delta(s)$*, we write* $s \xrightarrow{a}_n s'$*. In the first case, we speak of a probabilistic transition (and a probabilistic state* s*), in the second, we speak of a non-deterministic transition (and a non-deterministic state* s*). If a state* s *is non-deterministic and* $\delta(s) = \emptyset$*, we also call this state* final*. If a state* s *is non-deterministic and* $\delta(s)$ *a singleton set or empty set, or if* s *is probabilistic and* $\delta(s)$ *is the Dirac distribution, we also call* s deterministic*. A probabilistic transition system is fully probabilistic if all of its states are either probabilistic or deterministic.*

Here, $\mathcal{D}(\Omega)$ denotes the set of all discrete probability distributions on Ω. A discrete probability distribution on Ω is a function $\mu \colon \Omega \to [0, 1]$ such that $\{x \in \Omega \mid \mu(x) > 0\}$ (also called the support of μ, denoted $supp(\mu)$) is finite or countably infinite, and $\sum_{x \in \Omega} \mu(x) = 1$. We use $\mu(X)$ as short-hand for $\sum_{x \in X} \mu(x)$ for $X \subseteq \Omega$. For $x \in \Omega$, let μ_x^1 denote the Dirac distribution at x, i.e., the distribution with $\mu_x^1(x) = 1$.

Having this model in place, given an ADL program Π, we interpret Π as a transition system with a transition $\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as \rangle \xrightarrow{[p]}_n \langle m' \cdot ml', h', pp' \cdot ppl', r' \cdot rl' \rangle$ wherever $\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as \rangle \to_{\Pi, n, p} \langle m' \cdot ml, h', pp' \cdot ppl, r' \cdot rl, as' \rangle$.

Example 1: ADL transition system. Given an ADL program Π, an initial configuration $\langle ml, h, ppl, rl \rangle$, and an attacker system $(S_A, s_A^0, \emptyset, \delta_A)$ in $\{\mathcal{A}^\eta\}_{\eta \in \mathbb{N}} \in \mathcal{ADV}$, let $\mathrm{ADL}_{\Pi, \mathcal{A}, \langle ml, h, ppl, rl \rangle}$ be the probabilistic transition system $(\mathrm{S}, \langle ml, h, ppl, rl \rangle, \emptyset, \delta)$, where

- $\mathrm{S} = \mathcal{C} \uplus \mathcal{C}_{\text{final}} \uplus \mathcal{Q}_{final}$,
- $s_0 = \langle mk, h, ppl, rl, s_0^A \rangle$,
- $\delta(s) = \mu^1_{s',1}$ for s and s', according to the rules described in Figure 5 and Appendix A (observe that at most one rule applies to each state, and that each rule uniquely determines the follow-up state),
- $\delta(s) = \mu_{s,v_a,\mathcal{N}}$ where $\mu_{s,v_a,\mathcal{N}}(s\{r[v_a \mapsto n\}, 1) = \frac{1}{\mathcal{N}}$ for all $n \in \mathcal{N}$, if s matches the pre-condition of PROB in Figure 7,
- $\delta(s) = \mu_{s,\mathcal{A}}$, where $\mu_{s,\mathcal{A}}(s\{pp+1, r[res_{lo} \mapsto lo(res), res_{up} \mapsto up(res)], as'\}, n+1) = \Pr[\text{Exec}(T) = \langle (mid, r(v_a), \ldots, r(v_e)), as, \epsilon\rangle \to_A q_1 \to_A \cdots \to_A q_{n-1} \to_A \langle i, as', res\rangle]$ and $\mu_{s,\mathcal{A}}(\langle res\rangle) = \Pr[\text{Exec}(T) = \langle (mid, r(v_a), \ldots, r(v_e)), as, \epsilon\rangle \to_A q_1 \to_A \cdots \to_A q_{n-1} \to_A \langle i, as', res\rangle]$ for s according to the preconditions of ADVINV and ADVRET in the same figure. (Note that $s \in \mathcal{C}$.)
- $\delta(\langle u, h, as\rangle) = \mu_{as}$ where $\mu_{as} = \Pr[\text{Exec}(T) = \langle (), as, \epsilon\rangle \to_A q_1 \to_A \cdots \to_A q_{n-1} \to_A \langle i, as', res\rangle]$ and s according to the preconditions of ADVFIN. (Note that $s \in \mathcal{Q}_{final}$.)

$\diamond$

Definition 21 (Trace distribution (probabilistic))**.** *Given a fully probabilistic transition system* $T = (\mathrm{S}, s_0, A, \delta)$*, where* S *is finite or countably infinite and* $\epsilon \notin A$*, we define* $\Pr[s \xrightarrow{\alpha}_n s'] = p$ *iff.* $s \xrightarrow{[p]}_n s' \wedge \alpha = \epsilon$ *or* $s \xrightarrow{\alpha}_n s' \wedge p = 1$*. We define the outcome probability of an execution as follows:*

$$\Pr[\text{Exec}(T) = s_0 \xrightarrow{\alpha_1}_{n_1} \cdots \xrightarrow{\alpha_n}_{n_n} s_n] = \prod_{i=0}^{n-1} \Pr[s_i \xrightarrow{\alpha_{i+1}} s_{i+1}].$$

The probability to reach a certain state within n *steps is defined*

$$\Pr[T \downarrow_n s_m] = \sum \Pr\left[\text{Exec}(T) = s_0 \xrightarrow{\alpha_1}_{n_1} \cdots \xrightarrow{\alpha_m}_{n_m} s_m \text{ and } \sum_{i=1}^{m} n_i \leq n\right],$$

and the probability of a trace $(a_1, \ldots, a_n)$ *as*

$$\Pr[\mathit{Traces}(T) = (a_1, \ldots, a_n)] = \sum_{(\alpha_1,\ldots,\alpha_m)|_A = (a_1,\ldots,a_n)} \Pr[\text{Exec}(T) = s_0 \xrightarrow{\alpha_1}_{m_1} \cdots \xrightarrow{\alpha_m}_{m_m} s_m].$$

Notational conventions. Within $\Pr[\cdot]$, we abbreviate $s \xrightarrow{\alpha}_n s'$ with $s \xrightarrow{\alpha} s'$ if $n = 1$. We furthermore use α to indicate a transition that might be unlabelled, i.e., starts from a probabilistic state, in which case $\alpha = \epsilon$. We use a instead of α, to indicate that the parting state is indeed non-deterministic.

5.2 Threat model

In this work, we consider adversaries that are network parties or malicious apps that try to retrieve sensitive information from an honest app. We model the adversary as an external entity that cannot run any code within the program and that does not have access to the program's heap, but can read input to and control output from a given set of *malicious functions* $\mathcal{ADVR}$. The attacker is a probabilistic polynomial-time algorithm $\mathcal{A}$.

We represent the attacker $\mathcal{A}$ in ADL as an unlabelled probabilistic transition system. We assume that each state is a triple, where the first element solely contains the inputs and the last element contains the outputs. We use the relation $\to_A$ to denote transitions in the attacker's transition system

$$\textsc{Prob}\ \frac{m[pp] = \mathsf{rand}\ v_a \qquad n' \in \mathcal{N}}{s \xrightarrow{\left[\frac{1}{|\mathcal{N}|}\right]} s\{r[v_a \mapsto n']\}}$$

$$\textsc{AdvInv}\ \frac{m[pp] = \mathsf{invoke\text{-}static}\ \ v_a, \ldots, v_e, mid \qquad mid \in Mal_{static} \qquad res \in \mathcal{V} \qquad \Pr[\mathrm{Exec}(T) = \langle (mid, r(v_a), \ldots, r(v_e)), as, \epsilon\rangle \to_A q_1 \to_A \cdots \to_A q_{n-1} \to_A \langle i, as', res\rangle] = p}{s \xrightarrow{[p]}_{n+2} s\{pp+1, r[res_{lo} \mapsto lo(res), res_{up} \mapsto up(res)], as'\}}$$

$$\textsc{AdvRet}\ \frac{m[pp] = \mathsf{invoke\text{-}static}\ \ v_a, \ldots, v_e, mid \qquad mid \in Mal_{static} \qquad res \in \mathcal{ADVR} \qquad \Pr[\mathrm{Exec}(T) = \langle (mid, r(v_a), \ldots, r(v_e)), as, \epsilon\rangle \to_A q_1 \to_A \cdots \to_A q_{n-1} \to_A \langle i, as', res\rangle] = p}{s \xrightarrow{[p]}_{n+1} \langle res\rangle}$$

$$\textsc{AdvFin}\ \frac{\Pr[\mathrm{Exec}(T) = \langle (), as, \epsilon\rangle \to_A q_1 \to_A \cdots \to_A q_{n-1} \to_A \langle i, as', res\rangle] = p \qquad res \in \mathcal{ADVR}}{\langle u, h, as\rangle \xrightarrow{[p]}_{n+2} \langle res\rangle}$$

for $v_a, v_b, v_c, v_d, v_e \in \mathcal{X}$.

Figure 7: Inference rules extending ADL with probabilistic semantics and adversarial interaction.

in Figure 7. Many computation models can be expressed this way, including, but not restricted to, Turing machines.

Whenever a malicious function is invoked, $\mathcal{A}$ is executed with its previous state as and the arguments arg of the malicious function. The output of $\mathcal{A}$ is the new state as' and a response-message res in $\mathcal{V}$, or in the set $\mathcal{ADVR}$, which is distinct from $\mathcal{V}$. If $res \in \mathcal{V}$, it is interpreted as the function's output values, otherwise, i.e., if the response message is in $\mathcal{ADVR}$, the execution terminates with the adversarial output $res \in \mathcal{ADVR}$. Figure 7 precisely defines this behavior in the rule rInvoke-Adv.

Definition 22 (Attacker)**.** *An attacker $\mathcal{A}$ is a family of fully probabilistic transition systems (cf. Definition 20), indexed by a security parameter $\eta \in \mathbb{N}$. For each of theses probabilistic transition systems $(\mathrm{S}, s_0, \emptyset, \delta)$, the following holds: $\mathrm{S} = (F_{mal}, \mathcal{V}^*) \times \mathcal{Q} \times (\mathcal{V} \uplus \mathcal{ADVR} \uplus \{\epsilon\})$, where F_{mal} is a set of malicious functions, $\mathcal{Q}$ is the state-space of the adversary, and $\mathcal{ADVR}$ are the adversarial outputs, and each state s in $\mathcal{Q}_{final} := \{\langle i, q, res\rangle \mid res \in \mathcal{ADVR}\}$ is final, i.e., $\delta(s) = \emptyset$. Furthermore (for simplicity), we assume that every step has computation time 1, i.e., $\delta(s)(s', n) = 0$ for all $n \neq 1$.*

We define the adversarial computation model as a set of attackers, which we call $\mathcal{ADV}$.

Definition 23. *There is a subset of $\mathcal{MID}_\Pi$ which the adversary controls, called Mal. No element of Mal is in the range of any lookup-tables lookup-virtual, lookup-static, lookup-direct, lookup-super. Mal is partitioned into the sets $Mal_{virtual}$, Mal_{static}, Mal_{super}, and Mal_{direct}.*

We obtain the following definition for executing ADL in the presence of an adversary.

Definition 24 (ADL Execution with Adv)**.** *Given an ADL program Π, an adversary A, and an initial configuration $\langle ml, h, ppl, rl\rangle$, the probability that the interaction between Π on $\langle ml, h, ppl, rl\rangle$ and $\mathcal{A}$ terminates within n steps and results in x is defined as*

$$\Pr[\langle \Pi\, \langle ml, h, ppl, rl\rangle \,\|\mathcal{A}\rangle \downarrow_n x] := \Pr[\mathrm{ADL}_{\Pi, \mathcal{A}, \langle ml, h, ppl, rl\rangle} \downarrow_n \langle x\rangle].$$

where we write (given $\mathcal{A}$'s initial state as) $\Pi\, \langle ml, h, ppl, rl\rangle$ for the program Π with initial state $\langle ml, h, ppl, rl, as\rangle$.

5.3 Calls to crypto APIs

The most direct way of modelling the crypto-API would take the symbolic model as a starting point, defining static functions without any side-effect, one of which can be used to generate nonces or keys,

while the others apply constructors or destructors. However, as our results shall be applicable to Android Apps, we take existing crypto-APIs as the starting point, since the symbolic model is meant to abstract existing crypto-APIs, rather than demand a "matching implementation". To this end, it is worth having a look at how Java's Standard Cryptography Interface is typically used.

Listing 1: Excerpt from `https://www.owasp.org/index.php/Using_the_Java_Cryptographic_Extensions`

```
KeyGenerator keyGen = KeyGenerator.getInstance("AES");
keyGen.init(128);
SecretKey secretKey = keyGen.generateKey();
byte[] iv = new byte[128 / 8];
SecureRandom prng = new SecureRandom();
prng.nextBytes(iv);
Cipher c = Cipher.getInstance("AES/CBC/PKCS7PADDING");
c.init(Cipher.ENCRYPT_MODE,
       secretKey,
       new IvParameterSpec(iv));
byte[] byteCipherText = c.doFinal(byteDataToEncrypt);
```

We observe that here, `c` is a function object used for encryption and decryption, which is initialised with the cypher to be used, mode of operation and additional parameters. Consequently, when the actual computations are performed in the last line, the fields of this object carry information important to its operation, not only the arguments supplied. Hence, the cryptographic abstraction of this library depends not only on the method called (`doFinal` in this case), but also on the object in the heap. Moreover, the question *which* cryptographic abstraction is chosen, e.g., encryption, decryption, authenticated encryption, etc., depends on object in the heap.

Furthermore, these functions are typically `final`, meaning that they cannot be overwritten by sub classes. This is (most likely) compiled to a direct call of said function, implemented in the instruction invoke-direct-range (see Figure 6). From the point of view of this paper, invokevirtual, invokedirect, invoke-super and their -range counterparts behave very similarly, which is why we concentrate on the case relevant for Java's standard cryptographic interface, i.e., invoke-direct-range . But it is worth noting, that static function calls (invoke-static) are not used.

We can also observe, that the randomness used for encryption (`iv`) is generated using a PRNG, but user-supplied, while the `generateKey` method choses a key without exposing the randomness used. But even for keys, the randomness can be supplied by the user, as demonstrated by the following (insecure) example.

Listing 2: Excerpt from `https://gist.github.com/bricef/2436364`

```
static String IV = "AAAAAAAAAAAAAAAA";
static String plaintext = "test text 123\0\0\0"; /*Note null padding*/
static String encryptionKey = "0123456789abcdef";

Cipher cipher = Cipher.getInstance("AES/CBC/NoPadding", "SunJCE");
SecretKeySpec key = new SecretKeySpec(encryptionKey.getBytes("UTF-8"), "AES");
cipher.init(Cipher.ENCRYPT_MODE, key,new IvParameterSpec(IV.getBytes("UTF-8")));
return cipher.doFinal(plainText.getBytes("UTF-8"));
```

On the other hand, it is also possible to let the `Cipher` object pick the initialisation vector itself (and request it via a method call in order to, e.g., attach it to the ciphertext). This is implemented in Java via overloading, i.e., the method id is known at compile time based on the number and types of arguments.

To summarize these observations:

- Cryptographic output may depend on values in the heap.
- Which cryptographic abstraction is appropriate may depend on values in the heap, too.
- Cryptographic calls are calls to non-static functions.

- Cryptographic APIs provide interfaces for random number generation and key-generation with or without explicit randomness.

We can validate these observations on the open source library bouncycastle[4], which is very popular in mobile applications, as it is fairly lightweight. The interface it provides also permits user-supplied randomness, while at the same time retaining state within the `cipher` object.

Listing 3: Excerpt from `https://www.bouncycastle.org/specifications.html`

```
BlockCipher engine = new DESEngine();
BufferedBlockCipher cipher = new PaddedBlockCipher(new CBCCipher(engine));
byte[] key = keyString.getBytes();
byte[] input = inputString.getBytes();
cipher.init(true, new KeyParameter(key));
byte[] cipherText = new byte[cipher.getOutputSize(input.length)];
int outputLen = cipher.processBytes(input, 0, input.length, cipherText, 0);
cipher.doFinal(cipherText, outputLen);
```

Note that it is also possible to use bouncycastle as a security provider for Java's standard cryptographic interface.

Since randomness might be implicit, or explicit, depending on how a method is called, rather than mapping methods to (the evaluation of) a constructor or destructor on its arguments, we chose to generalize this to arbitrary combinations of constructors, destructors and randomness generation. Symbolic operations (see Definition 5) nicely capture the concept. Hence, a library specification maps a method and a predicate on the object's current state (i.e., a subset of $\mathcal{O}$) to a symbolic operation. The input to the symbolic operation contains the method's arguments, as well as a defined part of the heap (cf. Definition 39).

Definition 25 (Library specification)**.** *A library specification is an efficiently computable partial function* $libSpec : \mathcal{MID} \times \mathcal{O} \cup \mathcal{UNOP} \cup \mathcal{BINOP} \cup \mathcal{RELOP} \rightharpoonup \mathrm{SO}$ *defined at least on* $\mathcal{UNOP} \cup \mathcal{BINOP} \cup \mathcal{RELOP}$*.*

Note that operations on bitstrings are also specified by *libSpec*. In Corollary 3 we show that via symbolic operations, destructors derived from these operations can also be added as destructors (together with potentially some fresh nonces) to the symbolic model in a computationally sound way. As we will see later, this allows us to retain more precision, as bitstrings obtained from the crypto-API, e.g., via decryption, can still be treated as bitstrings.

In order to be able to use an asymptotic security definition, we define uniform families of ADL programs as programs generated from a security parameter.

Definition 26 (Uniform families of ADL programs)**.** *Let* Π *be an algorithm that, given a security parameter* η*, outputs an ADL program. We denote the output of this program* Π^η *and call the set of outputs of this program a* uniform family of ADL programs.

Next, we define initial configurations for families (indexed by a security parameter) of transition systems as states that are valid initial configuration for all security parameters.

Definition 27 (Initial configuration)**.** *Given a family* $(T_\eta)_\eta$ *of transition systems. We say that a state* s *is an* initial configuration *for* $(T_\eta)_\eta$ *if it is a valid initial configuration for all* T_η *in this family. Analogously, we say that a configuration* s *is an initial configuration for a uniform family of ADL programs if for all* η *this* s *is a valid initial configuration for* Π^η*.*

Real-world cryptographic libraries require each cipher before use to be initialised with the key-length, e.g., the class methods `getInstance` in `javax.crypto.Cipher` is called with a string specifying cipher, mode of operation, key-length and possibly more (see, e.g., Listing 2). Even though the key-length cannot be arbitrarily large (in fact, the choice is quite limited), we think this comes reasonably

[4] http://www.bouncycastle.org

close to how asymptotic security is achieved in real life: when standardisation bodies or security experts advise developers to chose larger keys, some constant in the existing source code is adapted to change the parameters to the security library, e.g., in a preprocessor step. We are aware that this requirement is not met by most real-world cryptographic libraries. Still, we consider this gap between the actually deployed programs and the theoretical result to be significantly smaller than in previous results for actual programming languages [6, 14, 15].

5.4 Indistinguishability of two ADL programs

Similar to CoSP, we define indistinguishability of ADL programs using tic-indistinguishability [12].

Definition 28 (Indistinguishability (ADL))**.** *We call two uniform families of ADL programs* $\Pi_1 = \{\Pi_1^\eta\}_{\eta\in\mathbb{N}}$ *and* $\Pi_2 = \{\Pi_2^\eta\}_{\eta\in\mathbb{N}}$ *with initial configuration* $s_1 = \langle ml_1, h_1, ppl_1, rl_1\rangle$ *and* $s_2 = \langle ml_2, h_2, ppl_2, rl_2\rangle$ computationally indistinguishable *for a (not necessarily uniform) family of attackers* $\mathcal{A} = \{\mathcal{A}^\eta\}_{\eta\in\mathbb{N}}$ $\mathcal{A}$ *(written* $\Pi_1\langle s_1\rangle \approx_c^{\mathrm{ADL}^{\mathcal{A}}} \Pi_2\langle s_2\rangle$*) if for all polynomials* p*, there is a negligible function* μ *such that for all* $a, b \in \{0, 1\}$ *with* $a \neq b$*,*

$$\Pr[\langle \Pi_1^\eta\langle s_1\rangle \| \mathcal{A}^\eta\rangle \downarrow_{p(\eta)} a] + \Pr[\langle \Pi_2^\eta\langle s_2\rangle \| \mathcal{A}^\eta\rangle \downarrow_{p(\eta)} b] \leq 1 + \mu(\eta).$$

We call $\Pi_1\langle s_1\rangle$ *and* $\Pi_2\langle s_2\rangle$ computationally indistinguishable ($\Pi_1\langle s_1\rangle \approx_c^{\mathrm{ADL}} \Pi_2\langle s_2\rangle$) *if we have* $\Pi_1\langle s_1\rangle \approx_c^{\mathrm{ADL}^{\mathcal{A}}} \Pi_2\langle s_2\rangle$ *for all machines* $\mathcal{A}$*.*

This notion indistinguishability gives rise to notion of non-interference, when $\Pi_1 = \Pi_2$.

6 Symbolic Dalvik Bytecode

In this section, we define a symbolic variant of ADL, which uses symbolic terms instead of cryptographic values. We sometimes abbreviate this symbolic variant as ADLs. We will show in the next section that it suffices to analyze a (symbolic) ADLs program in order to prove the corresponding (cryptographic) ADL program secure, provided that the cryptographic operations used in that program are computationally sound, i.e., provided that they have a computationally sound symbolic model in CoSPin the sense of Definition 15). ADLs' semantics precisely corresponds to the semantics of ADL, except for the treatment of cryptographic operations. As a consequence, existing automated analysis tools can be conveniently extended to ADLs and, thus, accurately cope with cryptographic operations, since this extension only requires semantic adaptations precisely for those cases where cryptographic behavior needs to be captured. ADLs is parametric in the symbolic model of the considered cryptographic operations and, hence, benefits from the rich set of cryptographic primitives that are already supported by computational soundness results in CoSP, such as encryption and signatures [6–8, 10] or zero-knowledge proofs [9, 11].

6.1 Embeddable symbolic CoSP models

We first define sufficient conditions under which a given CoSP model can be embedded into ADL. To this end, we require that values in $\mathcal{N}$ can be embedded into CoSP, i.e., there needs to be an injective function from $\mathcal{N}$ into $\mathbf{T}$, so that, e.g., a register value can be encrypted. As in previous embeddings [6, 14], we require that the symbolic model includes an equality operation *equals*.

Definition 29 (ADL-embeddable symbolic model)**.** *A symbolic model* $\mathbf{M} = (\mathbf{C}, \mathbf{N}, \mathbf{T}, \mathbf{D})$ *(see Definition 2) is* ADL-embeddable *if*

- *for* $S = \mathcal{V} \cup \mathcal{H} \cup Mal_{static} \cup \mathcal{UNOP} \cup \mathcal{BINOP} \cup \mathcal{RELOP} \cup \mathcal{CONV} \cup \{0, 1\}$*, there exists an injective function* $\iota : S \to \mathbf{T}$ *from* S *to terms consisting only of constructors, such that all*

$n \in S$ can be distinguished using deconstructors, i.e., for each $n \in S$, there are deconstructors $D_1, \ldots, D_m \in \mathbf{D}$ such that for any $n' \in S$,

$$D_1 \circ \cdots \circ D_m \circ \iota(n') = \begin{cases} \iota(n') & \text{if } n' = n \\ \bot & \text{if } n' \neq n. \end{cases}$$

We assume the inverse ι^{-1} on the range of ι to be efficiently computable.

- *The destructor distinguishing $0 \in S$ is called iszero.*
- *There is a destructor $equals/2 \in \mathbf{D}$ such that for all $x, y \in \mathbf{T}$, $x \neq y$, $equals(x, x) = x$ and $equals(x, y) = \bot$.*
- *There are $pair/2 \in \mathbf{C}$ and $fst/1, snd/1 \in \mathbf{D}$ such that for all $, y \in \mathbf{T}$, $fst(pair(x, y)) = x$ and $snd(pair(x, y)) = y$.*
- *For each n-ary operation $op \in \mathcal{UNOP} \cup \mathcal{BINOP} \cup \mathcal{RELOP}$, there is a destructor op_s such that for all $a_1, \ldots, a_n \in \mathcal{V}$, $op_s(\iota(a_1), \ldots, \iota(a_n)) = \iota(\underline{op}(\iota^{-1}(a_1), \ldots, \iota^{-1}(a_n)))$.*

We abbreviate $\widehat{n} := \iota(n)$, and call $\widehat{n}$ the symbolic representation of n. We lifts this notion to sets: $\forall N \subseteq \mathcal{N}.\widehat{N} = \{\widehat{n} \mid n \in N\}$.

With Corollary 2, all recent CoSP-results from the literature [6–11] satisfy these requirements. We show that the first condition of Definition 29 can easily be satisfied for ADL values, if we instantiate the (underspecified) set $\mathcal{V}$ with bitstrings.

Example 2: ADL-embeddable symbolic model for bitstrings and heaps.

Let $\mathcal{V} = \{0,1\}^*$, i.e., numerical values, locations (and void) are expressed via bitstrings, but are distinguishable, i.e., via a tagging convention. We show how an embedding of bitstring can be achieved with a symbolic model $\mathbf{M} = (\mathbf{C}, \mathbf{N}, \mathbf{T}, \mathbf{D})$ which includes

- $equals/2 \in \mathbf{D}$ such that for all $x, y \in \mathbf{T}$, $x \neq y$, $equals(x, x) = x$ and $equals(x, y) = \bot$, and
- $string_0/1, string_1/1, emp/0 \in \mathbf{C}$, $unstring_0/1, unstring_1/1 \in \mathbf{D}$ such that for all $x \in \mathbf{T}$

$$unstring_0(string_0(x)) = x \text{ and } unstring_1(string_1(x)) = x,$$

and in all other cases $unstring_0(x) = \bot$ and $unstring_1(x) = \bot$.

In this case, $\iota_{\mathcal{V}} : \{0,1\}^* \to \mathbf{T}$ can be defined as follows:

$$\iota(b_l \ldots b_0) := string_{b_l}(string_{b_{l-1}}(\ldots(string_{b_0}(emp))\ldots)).$$

The symbols in $Mal_{static} \cup \mathcal{UNOP} \cup \mathcal{BINOP} \cup \mathcal{RELOP} \cup \mathcal{CONV} \cup \{0,1\}$ can be expressed using the same method via an embedding into bitstrings, or even simpler, by having one constructor and one destructor per element, as these sets are a-priori fixed. One could make the same argument for the heap, but as symbolic abstractions may produce updates to the heap, i.e., "partial" heaps that overwrite the heap where they are defined (cf. Section 6.2), and symbolic verification tools do not support arbitrary destructors, we propose a modelling as a list of pairs of argument and function value. As the heap can contain objects, as well as arrays, these need to be embedded as well. ◇

For illustration consider the bitstring `01101`. Its *symbolic representation* is the term

$$t := \widehat{01101} = string_0(string_1(string_1(string_0(string_1(emp))))).$$

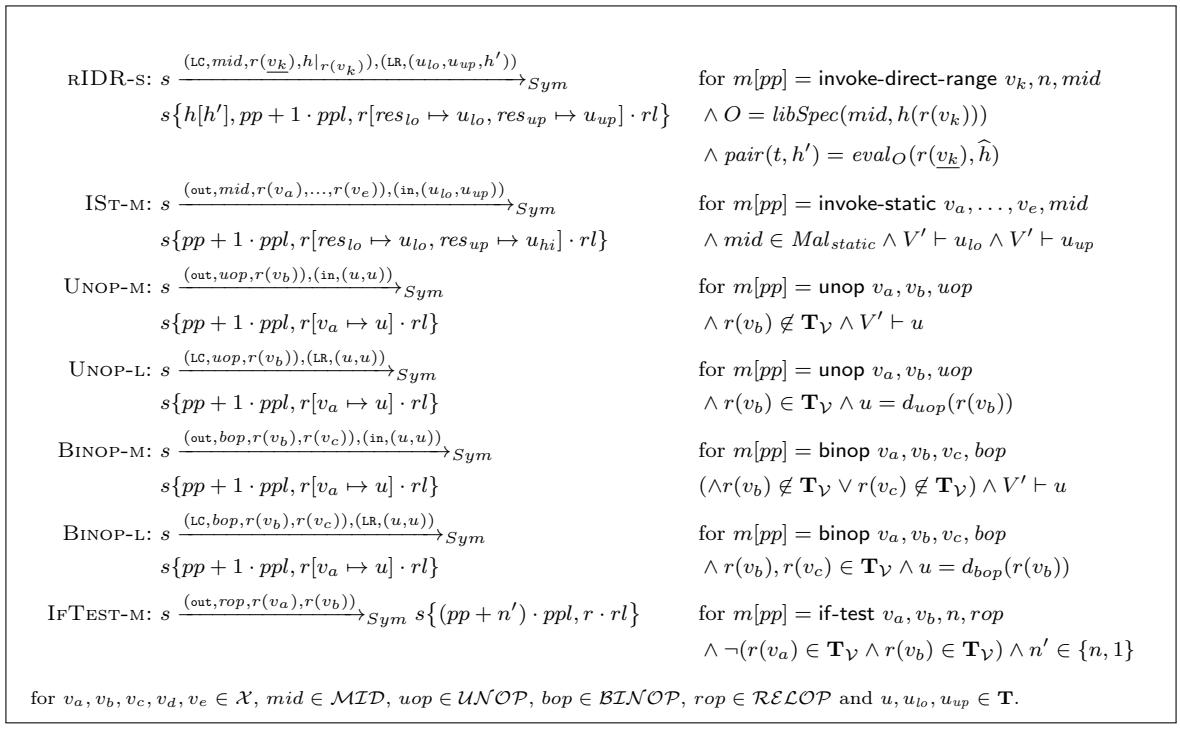
rIDR-s: $s \xrightarrow{(\texttt{LC}, mid, r(\underline{v_k}), h|_{r(v_k)}), (\texttt{LR}, (u_{lo}, u_{up}, h'))} Sym$ $s\{h[h'], pp+1\cdot ppl, r[res_{lo}\mapsto u_{lo}, res_{up}\mapsto u_{up}]\cdot rl\}$ for $m[pp]=$ invoke-direct-range v_k, n, mid $\wedge O = libSpec(mid, h(r(v_k)))$ $\wedge pair(t,h') = eval_O(r(\underline{v_k}), \widehat{h})$
ISt-m: $s \xrightarrow{(\texttt{out}, mid, r(v_a), \ldots, r(v_e)), (\texttt{in}, (u_{lo}, u_{up}))} Sym$ $s\{pp+1\cdot ppl, r[res_{lo}\mapsto u_{lo}, res_{up}\mapsto u_{hi}]\cdot rl\}$ for $m[pp]=$ invoke-static $v_a, \ldots, v_e, mid$ $\wedge mid \in Mal_{static} \wedge V' \vdash u_{lo} \wedge V' \vdash u_{up}$
Unop-m: $s \xrightarrow{(\texttt{out}, uop, r(v_b)), (\texttt{in}, (u,u))} Sym$ $s\{pp+1\cdot ppl, r[v_a\mapsto u]\cdot rl\}$ for $m[pp]=$ unop v_a, v_b, uop $\wedge r(v_b)\notin \mathbf{T}_{\mathcal{V}} \wedge V' \vdash u$
Unop-l: $s \xrightarrow{(\texttt{LC}, uop, r(v_b)), (\texttt{LR}, (u,u))} Sym$ $s\{pp+1\cdot ppl, r[v_a\mapsto u]\cdot rl\}$ for $m[pp]=$ unop v_a, v_b, uop $\wedge r(v_b)\in \mathbf{T}_{\mathcal{V}} \wedge u = d_{uop}(r(v_b))$
Binop-m: $s \xrightarrow{(\texttt{out}, bop, r(v_b), r(v_c)), (\texttt{in}, (u,u))} Sym$ $s\{pp+1\cdot ppl, r[v_a\mapsto u]\cdot rl\}$ for $m[pp]=$ binop v_a, v_b, v_c, bop $(\wedge r(v_b)\notin \mathbf{T}_{\mathcal{V}} \vee r(v_c)\notin \mathbf{T}_{\mathcal{V}}) \wedge V' \vdash u$
Binop-l: $s \xrightarrow{(\texttt{LC}, bop, r(v_b), r(v_c)), (\texttt{LR}, (u,u))} Sym$ $s\{pp+1\cdot ppl, r[v_a\mapsto u]\cdot rl\}$ for $m[pp]=$ binop v_a, v_b, v_c, bop $\wedge r(v_b), r(v_c)\in \mathbf{T}_{\mathcal{V}} \wedge u = d_{bop}(r(v_b))$
IfTest-m: $s \xrightarrow{(\texttt{out}, rop, r(v_a), r(v_b))} Sym\ s\{(pp+n')\cdot ppl, r\cdot rl\}$ for $m[pp]=$ if-test v_a, v_b, n, rop $\wedge \neg(r(v_a)\in \mathbf{T}_{\mathcal{V}} \wedge r(v_b)\in \mathbf{T}_{\mathcal{V}}) \wedge n'\in\{n,1\}$
for $v_a, v_b, v_c, v_d, v_e \in \mathcal{X}$, $mid \in \mathcal{MID}$, $uop \in \mathcal{UNOP}$, $bop \in \mathcal{BINOP}$, $rop \in \mathcal{RELOP}$ and $u, u_{lo}, u_{up} \in \mathbf{T}$.

Figure 8: Modified inference rules of the symbolic execution relation $\rightarrow$ -Sym, where $s = \langle ml, h, pp\cdot ppl, r\cdot rl\rangle$. (Variable ranges as in Figure 6.)

If a term t represents the bitstring `01101`, the following equality test evaluates to emp:

$$
\begin{aligned}
&equals(unstring_1(unstring_0(unstring_1(unstring_1(unstring_0(t))))), emp)\\
&= equals(unstring_1(unstring_0(unstring_1(unstring_1(unstring_0(string_0(string_1(string_1(string_0(string_1(emp))))))))), emp)\\
&= equals(unstring_1(unstring_0(unstring_1(unstring_1(string_1(string_1(string_0(string_1(emp))))))), emp)\\
&= equals(unstring_1(unstring_0(unstring_1(string_1(string_0(string_1(emp)))))), emp)\\
&= equals(unstring_1(unstring_0(string_0(string_1(emp)))), emp)\\
&= equals(unstring_1(string_1(emp)), emp)\\
&= equals(emp, emp)\\
&= emp
\end{aligned}
$$

In Appendix B, we show that any computationally sound symbolic model that contains symbolic bitstrings and equality can be extended by destructors implementing any polynomial-time computable function on bitstrings, allowing us to perform, e.g., binary operations like XOR on the symbolic representation of bitstrings.

6.2 Semantics of symbolic ADL

The semantical domains of the symbolic variant of ADL coincide with the semantical domains of ADL except for the set of values $\mathcal{V}$ and the set of (intermediate and final) states. The set of values is the union of symbolic terms and values as defined previously.

We add to each state in the inference rules the list of terms output so far, called *the symbolic view*. This symbolic view corresponds to the symbolic view in the symbolic execution of a CoSP protocol (Definition 6), and is used to track the information the adversary can use to deduce messages.

Definition 30 (Symbolic semantical domains)**.** *The* symbolic semantical domains *of ADL programs are defined as in Definition 16, with the exception of*

$$\mathcal{V}^{\mathbf{T}} = \mathcal{N} \cup \mathcal{L} \cup \{\text{void}\} \cup \mathbf{T} \qquad \textit{values}$$
$$\mathcal{C}^{\mathbf{T}} = \mathcal{H} \times \mathbb{N} \times \mathcal{R} \times \text{SViews} \qquad \textit{intermediate states}$$
$$\mathcal{C}^{\mathbf{T}}_{final} = \mathcal{V} \times \mathcal{H} \times \text{SViews} \qquad \textit{final states}$$

Transition between states are now annotated with the information sent to or received from the adversary. As ADL is sequential, every output to the adversary is followed by the adversary's input. Calls to malicious functions reveal register values to the adversary, which can be terms, locations, or numerical values. Terms are added to the adversarial knowledge as they are; locations and values are translated by means of the embedding. Calls to the crypto-API immediately yield a term by applying a constructor or destructor to the term arguments, or their symbolic representation. As elaborated in the introduction, applying binary operations (e.g., XOR) to symbolic terms usually invalidates computational soundness results. We over-approximate this by treating symbolic terms as a blackbox and let the outcome of any operation involving a symbolic term be decided by the adversary. As a side effect, we let the adversary learn the operands. As binary operations and tests on bits are only defined on inputs in $\mathcal{N}$, rules involving these remain unaltered and are complemented by the rules in Figure 8. For brevity, we let V' denote the attacker's view, which is part of the successor state, and write $V' \vdash m$, if the adversary can deduce a term m from a view V, i.e., there is a symbolic operation O such that $O(Out(V)) = m$. Observe that for IfTest-m, the adversary can decide the outcome without learning the operands. This highlights the non-determinism inherent to the symbolic semantics. We thus define the probabilistic transition system $\text{ADLs}_{\Pi,\langle ml,h,ppl,rl\rangle}$ similar to Example 1, but with these modifications.

With the symbolic semantics in place, we can finally define the notion of symbolic equivalence between two programs. We first specify what constitutes a symbolic view, as Definition 21 applies only to fully probabilistic transition systems.

Definition 31 (Symbolic view(probabilistic transition system))**.** *Given a probabilistic (but not necessarily fully probabilistic) transition system T with initial state s_0, the set of symbolic views of T is defined:*

$$\text{SViews}(T) = \left\{ (\alpha_1, \ldots, \alpha_m)|_{Views} \middle| s_0 \xrightarrow{\alpha_1}_{n_1} \cdots \xrightarrow{\alpha_m}_{n_m} s_m \right\}$$

As in CoSP, we use the equivalence relation on traces introduced in Definition 7.

Definition 32 (Symbolic equivalence)**.** *Let $\mathbf{M}$ be a symbolic model. Two probabilistic transition systems T_1 and T_2 are* symbolically equivalent *($T_1 \approx_s^{\text{SS}} T_2$) if $\text{SViews}(T_1) \sim \text{SViews}(T_2)$ w.r.t. $\mathbf{M}$. Two uniform families of ADL programs $\Pi_1 = \{\Pi_1^\eta\}_{\eta\in\mathbb{N}}$ and $\Pi_2 = \{\Pi_2^\eta\}_{\eta\in\mathbb{N}}$ and initial configurations $s_1 = \langle ml_1, h_1, ppl_1, rl_1\rangle$ and $s_2 = \langle ml_2, h_2, ppl_2, rl_2\rangle$ are* symbolically equivalent *($\Pi_1\langle s_1\rangle \approx_s^{\text{ADL}} \Pi_2\langle s_2\rangle$) iff, for each η, the probabilistic transition systems $T_1 = \text{ADLs}_{\Pi_1,s_1}$ and $T_2 = \text{ADLs}_{\Pi_2,s_2}$ are symbolically equivalent.*

7 Computational Soundness

Establishing a computational soundness proof for ADL requires a clear separation between honest program parts and cryptographic API calls with their corresponding augmented adversarial symbolic capabilities. To achieve this, we characterize this partitioning by introducing the concept of a *split-state form* of an operational semantics, and we subsequently show that it can be naturally used to represent ADL. This section, presents the essence of our proof. For the full proofs and detailed definition, we refer to the technical report [16].

7.1 Split-state semantics

The split-state form partitions the original semantics into three components, parallely executed asynchronously: (*i*) all steps that belong to computing cryptographic operations (called the *crypto-API semantics*), (*ii*) all steps that belong to computing the malicious functions (called the *attacker semantics*), and (*iii*) all steps that belong to the rest of the program (called the *honest-program semantics*). The overall operational semantics is a composition of each of these sub-semantics, i.e., the state space is the Cartesian product of the sub-states with additional information on which entity is currently running, defined in terms of asynchronous parallel composition. In a split-state form, all three entities can synchronize through the following sets of labels: libCall, libResp, out, and in. The first two model message passing between the honest program and the crypto-API; the latter two model message passing between the honest program and the attacker. There is no synchronization step between attacker and the crypto-API, as the attacker can only stop the execution with a transition of the form $(\texttt{final}, m)$ for some message m.

The overall operational semantics, called the *split-state semantics*, will then be defined as the composition of each of these three sub-semantics, i.e., the state space is the Cartesian product of the sub-states with additional information on which entity is currently running, defined in terms of asynchronous parallel composition. The benefit of this notion is that it makes the communication explicit that occurs between the actual program and the adversary. This communication is often fixed but arbitrary, i.e., in the case of black-box usage. Moreover, it separates the actual program from the cryptographic API calls, which is crucial for computational soundness proofs. Finally, the details of the crypto-API are irrelevant for program analyses, provided that the cryptographic operations are implemented securely.

Defining split-state. In the first step, we represent the respective program semantics of the adversary, the cryptographic library and the honest programs in terms of three (possibly different) labelled transition systems. In each of those, every state corresponds to a configuration of the program, the set of initial states corresponds to the set of initial configurations, and the transition relation is defined by the set of possible execution steps. Execution steps that involve any of the other entities are labelled either:

- $(\texttt{LC}, f, m)$ when the honest program makes a library call to the function named f with input m.
- $(\texttt{LR}, m)$ for the library's response
- $(\texttt{out}, m)$ when the honest program passes message m to the adversary, and
- $(\texttt{in}, m)$ in the opposite case.

We can describe our notion of split-state semantics using the well-studied notion of *asynchronous parallel composition* [17,18][5], which we review below. For readability, we define the transition function in terms of the transition relation it induces.

Definition 33 (Asynchronous parallel composition)**.** *The asynchronous parallel composition of two probabilistic transition systems* $(\mathrm{S}_1, s_1^0, A_1, \delta_1)$ *and* $(\mathrm{S}_2, s_2^0, A_2, \delta_2)$ *is the probabilistic transition system* $(\mathrm{S}_1 \times \mathrm{S}_2, (s_1^0, s_2^0), (A_1 \cap A_2), \delta)$*, where* δ *is defined such that:*

- $(s_1, s_2) \xrightarrow{a}_{n_1+n_2} (s_1', s_2')$ *iff* $s_1 \xrightarrow{a}_{n_1} s_1'$, $s_2 \xrightarrow{a}_{n_2} s_2'$ *and* $a \in A_1 \cap A_2$,
- $(s_1, s_2) \xrightarrow{a}_{n_1} (s_1', s_2)$ *iff* $s_1 \xrightarrow{a}_{n_1} s_1' \wedge a \notin A_2$,
- $(s_1, s_2) \xrightarrow{a}_{n_2} (s_1, s_2')$ *iff* $s_2 \xrightarrow{a}_{n_2} s_2' \wedge a \notin A_1$, *and*
- $(s_1, s_2) \xrightarrow{[p]}_{n} (s_1', s_2')$ *iff* $s_1 \xrightarrow{[p]}_{n} s_1' \wedge s_2' = s_2$ *or* $s_2 \xrightarrow{[p]}_{n} s_2' \wedge s_1' = s_1$.

[5] To be more precise: CSP-style parallel composition on common actions.

We extend this form of composition to an arbitrary number of transition systems by applying it to the first and second element, then applying the result to the third, and so forth. We simplify notation by flattening the resulting state space to $\mathrm{S}_1 \times \mathrm{S}_2 \cdots \times \mathrm{S}_n$ instead of $(\cdots(\mathrm{S}_1 \times \mathrm{S}_2)\cdots) \times \mathrm{S}_n$, modifying initial configurations and transition relation accordingly.

A split-state composition is an asynchronous parallel composition with the following restrictions: *a*) We specify which transitions can be performed in synchronisation. Assuming some set of function symbols F, some set of symbols for malicious functions F_{mal} including `finalCall`, and a domain for messages $\mathcal{D}$, we define:

$$\begin{aligned}
\mathrm{out} &:= \{((\texttt{out}, f, m_1, \ldots, m_n), 1, 1) \mid f \in F_{mal}, m_i \in \mathcal{D}\} \\
\mathrm{in} &:= \{((\texttt{in}, m), 1, 1) \mid m \in \mathcal{D}\}, \text{ and} \\
\mathrm{libCall} &:= \{((\texttt{LC}, f, m_1, \ldots, m_n), 1, 1) \mid f \in F, m_i \in \mathcal{D}\}, \\
\mathrm{libResp} &:= \{((\texttt{LR}, m), 1, 1) \mid m \in \mathcal{D}\}, \\
\mathrm{final} &= \{((\texttt{final}, m), 1, 1) \mid m \in \mathcal{D}\}.
\end{aligned}$$

The first two allow for message passing between the honest program and the cryptographic library, the latter two for message passing between honest program and attacker. There is no synchronising step between attacker and library. *b*) We enrich the three transition system's state with a boolean value that indicates whether the system is currently active. Only the active system can perform non-synchronising steps or decide to pass information to another system via a synchronising step – but in this case, it activates the other system and ceases to be active itself.

To simplify notation, we write $s \xrightarrow{A} s'$ if there exists $a \in A$ such that $s \xrightarrow{a} s'$.

Definition 34 (Split-state operational semantics)**.** *Given three transition systems* $(\mathrm{S}_H, s_H^0, A_H, \delta_H)$, $(\mathrm{S}_A, s_A^0, A_A, \delta_A)$, $(\mathrm{S}_L, s_L^0, A_L, \delta_L)$, *such that* $A_H \subseteq \mathrm{out} \cup \mathrm{in} \cup \mathrm{libCall} \cup \mathrm{libResp}$, $A_A \subseteq \mathrm{out} \cup \mathrm{in} \cup \mathrm{final}$, $A_L \subseteq \mathrm{libCall} \cup \mathrm{libResp}$, *we define the split-state composition as the asynchronous parallel composition of* $(\mathbb{B} \times \mathrm{S}_H, (\top, s_H^0), A_H, \delta'_H)$, $(\mathbb{B} \times \mathrm{S}_A, (\bot, s_A^0), A_A, \delta'_A)$, $(\mathbb{B} \times \mathrm{S}_L, (\bot, s_L^0), A_L, \delta'_L)$, *where*

$$\delta'_H(b,s) = \begin{cases} \mu \circ (((\top, s), n) \mapsto (s, n)) & \text{if } b = \top \text{ and } \delta_H(s) = \mu \\ ((\texttt{out}, \texttt{finalCall}), (\bot, s), 1) & \text{if } b = \top \text{ and } s \text{ final} \\ \delta_H(s)|^{\bot}_{\mathrm{libCall} \cup \mathrm{out}} & \text{if } b = \top \text{ and } s \text{ non-determ. and not final} \\ \delta_H(s)|^{\top}_{\mathrm{libResp} \cup \mathrm{in}} & \text{if } b = \bot \text{ and } s \text{ non-deterministic} \end{cases}$$

$$\delta'_A(b,s) = \begin{cases} \mu \circ (((\top, s), n) \mapsto (s, n)) & \text{if } b = \top \text{ and } \delta_A(s) = \mu \\ \delta_A(s)|^{\bot}_{\mathrm{in}} & \text{if } b = \top \text{ and } s \text{ non-deterministic} \\ \delta_A(s)|^{\top}_{\mathrm{out}} & \text{if } b = \bot \text{ and } s \text{ non-deterministic} \end{cases}$$

$$\delta'_L(b,s) = \begin{cases} \mu \circ (((\top, s), n) \mapsto (s, n)) & \text{if } b = \top \text{ and } \delta_L(s) = \mu \\ \delta_L(s)|^{\bot}_{\mathrm{libResp}} & \text{if } b = \top \text{ and } s \text{ non-deterministic} \\ \delta_L(s)|^{\top}_{\mathrm{libCall}} & \text{if } b = \bot \text{ and } s \text{ non-deterministic} \end{cases}$$

Here, $P|^b_S$ *is short-hand for* $\{(a, (b, s'), n) \mid (a, s', n) \in P \wedge a \in S\}$. *An operational semantics is a* split-state semantics*, if it can be expressed as a split-state composition of three transition systems.*

We call states of the honest program semantics configurations*, and write* $T(s)$ *for a split-state semantics* T *meaning that the initial state of the honest program semantics* T_H *is replaced by* s.

We stress that requiring an `out`-transition being always followed by a `in`-transition is without loss of generality. In case the original would solely send a message or solely receive a message, one of the transitions would only carry dummy information.

The benefit of this notion is that it makes the communication explicit that occurs between the actual program and the adversary. This communication is often fixed but arbitrary, i.e., in the case of

black-box usage. Moreover, it separates the actual program from the cryptographic API calls, which is crucial for computational soundness proofs. Finally, the details of the crypto-API are irrelevant for program analyses, provided that the cryptographic operations are implemented securely.

This notion makes these distinction more explicit and facilitates the proof, but note that the split-state semantics can be defined in a meaningless way: the library could be empty, and never be called, or return values from the adversary be wholly ignored. We will illustrate in the following section, however, that through small additions, the ADL semantics can be separated, providing *a*) for a generalized proof method, i.e., Lemma 11 and 13 apply to all split-state semantics, not only ADL, *b*) a communication model that makes the threat model and the assumptions on the library implementation obvious.

A split-state semantics is not necessarily a probabilistic transition systems, as the transition systems it is composed from might be non-deterministic. However, we define the probability of a certain outcome only for probabilistic split-state semantics, as in most applications for security, non-determinism in the honest program semantics stems from the modelling of concurrency and is usually resolved by specifying a scheduler. This non-determinism is typically conservatively resolved by assuming that the adversary controls the scheduling.

Definition 35 (Split-state execution)**.** *Let T_H, T_L, T_A be transition systems and assume that their split-state composition T is probabilistic.*

We write $\Pr[T(s_0) \downarrow_n x]$ *for the probability that the interaction between between honest program, attacker and library with initial state $s_0 \in \mathrm{S}^0$ results in x and terminates within n steps, as follows:*

$$\Pr[T(s_0) \downarrow_n x] = \sum_{(s_0 \xrightarrow{\alpha_1} \cdots s_k) \in supp(\mathrm{Exec}(T))} \Pr\left[\mathrm{Exec}(T) = s_0 \xrightarrow{\alpha_1}_{n_1} s_1 \cdots \xrightarrow{\alpha_k = (\mathtt{final}, x)}_{n_k} s_k \text{ and } k \leq n\right].$$

Definition 36 (Split-state indistinguishability)**.** *Let $T_{H,1}$, $T_{H,2}$, T_A, and T_L be families of transition systems indexed by a security parameter in $\mathbb{N}$. We call $T_{H,1}$ and $T_{H,2}$* computationally indistinguishable *for T_L and T_A and initial states $s^0_{H,1}$, $s^0_{H,2}$, s_L of $T_{H,1}$, $T_{H,2}$, T_L (respectively) in the sense of Definition 27 if for all polynomials p, there is a negligible function μ such that for all $a, b \in \{0,1\}$ with $a \neq b$,*

$$\begin{aligned}&\Pr[T_1^\eta((\top, s^0_{H,1}), (\bot, (0, z)), (\bot, s_L)) \downarrow_{p(\eta)} a] \\ &+ \Pr[T_2^\eta((\top, s^0_{H,2}), (\bot, (0, z)), (\bot, s_L)) \downarrow_{p(\eta)} b] \leq 1 + \mu(\eta),\end{aligned}$$

where T_1^η is the split-state composition of $T^\eta_{H,1}, T^\eta_L, T_A$, and T_2^η is the split-state composition of $T^\eta_{H,2}, T^\eta_L, T_A$ for security parameter η. In short, we write $T_{H,1}(s^0_{H,1}) \approx_c^{T_L(s_L), T_A} T_{H,2}(s^0_{H,2})$. Additionally, we call $T_{H,1}$ and $T_{H,2}$ computationally indistinguishable (written as $T_{H,1}(s^0_{H,1}) \approx_c^{T_L(s_L)} T_{H,2}(s^0_{H,2})$) *if $T_{H,1}(s^0_{H,1}) \approx_c^{T_L(s_L), T_A} T_{H,2}(s^0_{H,2})$ for all $T_A \in \mathcal{ADV}'$ and for initial states $s^0_{H,1}, s^0_{H,2}, s_L$ for $T_{H,1}, T_{H,2}, T_L$ (respectively).*

7.2 Split-state representation of ADL

We now define the condition necessary for splitting an ADL program into an honest program semantic and a crypto-API semantics. An ADL program Π is *pre-compliant* with a crypto-API specification if: (i) the crypto-API only reads the slice of the heap belonging to its object, i.e., only objects reachable from locations stored in its own member variables, (ii) the crypto-API never invokes the adversary, and (iii) only the crypto-API and the adversary make use of the rand-instruction.

Definition 37 (Reachable locations)**.** *The set of reachable locations from location $l \in \mathcal{L}$ in heap $h \in \mathcal{H}$*

$$\textsc{LibCall}\colon s \xrightarrow{(\texttt{LC}, mid, r(v_k), \ldots, r(v_{k+n-1}), h|_{r(v_k)})}_H (\texttt{wait}, s) \qquad \begin{array}{l}\text{for } (mid, h(r(v_k))) \in dom(libSpec)\\ m[pp] = \textsf{invoke-direct-range } v_k, n, mid\end{array}$$

$$\textsc{LibResponse}\colon (\texttt{wait}, s) \xrightarrow{(\texttt{LR}, (m_{lo}, m_{up}, h'))}_H s \left\{ \begin{array}{l} h[h'], pp + 1 \cdot ppl, \\ r\,[res_{lo} \mapsto m_{lo}, res_{up} \mapsto m_{up}] \cdot rl \end{array} \right\} \qquad \begin{array}{l}\text{for } (mid, h(r(v_k))) \in dom(libSpec)\\ m[pp] = \textsf{invoke-direct-range } v_k, n, mid\end{array}$$

$$\textsc{LeakMsg}\colon s \xrightarrow{(\texttt{out}, mid, r(v_a), \ldots, r(v_e))}_H (\texttt{wait}, s) \qquad \begin{array}{l}\text{for } mid \in Mal_{static}\\ m[pp] = \textsf{invoke-static } v_a, \ldots, v_e, mid\end{array}$$

$$\textsc{ReceiveMsg}\colon (\texttt{wait}, s) \xrightarrow{(\texttt{in}, m_{lo}, m_{up})}_H s \left\{ pp + 1 \cdot ppl, r \begin{bmatrix} res_{lo} \mapsto m_{lo}, \\ res_{up} \mapsto m_{up} \end{bmatrix} \cdot rl \right\} \qquad \begin{array}{l}\text{for } mid \in Mal_{static}\\ m[pp] = \textsf{invoke-static } v_a, \ldots, v_e, mid\end{array}$$

$$\textsc{FinalCall}: \langle u, h\rangle \xrightarrow{(\texttt{out}, \texttt{finalCall})}_H \langle u, h\rangle$$

for $v_a, v_b, v_c, v_d, v_e, v_k \in \mathcal{X}$, $n \in \mathbb{N}$, $mid \in \mathcal{MID}$, and $m_{lo}, m_{up} \in \mathcal{V}$

Figure 9: ADL split-state representation, honest program semantics, $s = \langle ml, h, ppl, r \cdot rl\rangle$

$$\textsc{LLibCall}: \langle (), \emptyset, (), (), \_\rangle \xrightarrow{(\texttt{LC}, mid, r(\underline{v}), h)}_L \langle (m), h, (0), (defReg([r(\underline{v})])), \_\rangle \qquad \begin{array}{l}\text{for } r(\underline{v}) := r(v_k), \ldots, r(v_{k+n-1}),\\ m := lookup\text{-}direct_\Pi(mid, h(r(v_k))).\textsf{class}\end{array}$$

$$\textsc{LLibRetVoid}: \langle (m), h, (pp), (r), \_\rangle \xrightarrow{(\texttt{LR}, (\text{void}, \text{void}, h))}_L \langle (), \emptyset, (), (), \_\rangle \qquad \text{for } m[pp] = \textsf{return-void}$$

$$\textsc{LLibRet}: \langle (m), h, (pp), (r), \_\rangle \xrightarrow{(\texttt{LR}, (lo(r(v_a)), up(r(v_a)), h))}_L \langle (), \emptyset, (), (), \_\rangle \qquad \text{for } m[pp] = \textsf{return } \ v_a$$

for $v_a, v_b, v_k \in \mathcal{X}$, $n \in \mathbb{N}$, $mid \in \mathcal{MID}$.

Figure 10: ADL split-state representation, library program semantics, _ is an arbitrary adversarial state.

is defined recursively:

$$lreachable_h(l) = \begin{cases} l & \text{if } h(l) \in \mathcal{AR} \\ \emptyset & \text{if } h(l) \text{ undefined} \\ \bigcup_{f|F(f)\in\mathcal{L}} lreachable_h(l) & \text{if } h(l) = (c, F) \in \mathcal{O} \end{cases}$$

Definition 38 (Heap slice)**.** *Let $h \in \mathcal{H} = \mathcal{L} \rightharpoonup (\mathcal{O} \cup \mathcal{AR})$ and $l \in \mathcal{L}$. We define l's slice of h as a partial function on $\mathcal{L}$ as follows:*

$$h|_l(l') = h(l') \quad \text{if } l' \in lreachable_h(l)$$

Definition 39 (ADL pre-compliance)**.** *An ADL program Π is pre-compliant with a library specification libSpec, if for each η for the Crypto-API semantics in $\text{ADL}^{SS}_{\Pi^\eta, \mathcal{A}^\eta, s_0}$ (see Definition 40) the following holds:*

- *the Crypto-API semantics only reads the slice of the heap belonging to its object, i.e., for all $mid \in \mathcal{MID}$, $h_1, h_2 \in \mathcal{H}$, $v_k, \ldots, v_{k+n-1} \in \mathcal{V}$, $as \in \mathcal{Q}$, $u \in \mathcal{V}$, such that $(mid, h(r(v_k))) \in libSpec$ and $m = lookup\text{-}direct_\Pi(mid, h(r(v_k))).\textsf{class}$,*

$$\Pr[\langle (m), h_1[h|_{v_k}], (0), (defReg([r_k, \ldots, r_{k+n-1}])), as\rangle \xrightarrow{\alpha_1} \cdots \xrightarrow{\alpha_n} \langle u, h_1[h|_{v_k}], as_n\rangle] =$$
$$\Pr[\langle (m), h_2[h|_{v_k}], (0), (defReg([r_k, \ldots, r_{k+n-1}])), as\rangle \xrightarrow{\alpha_1} \cdots \xrightarrow{\alpha_n} \langle u, h_2[h|_{v_k}], as_n\rangle].$$

- *the Crypto-API semantics never invokes the adversary, i.e., for all $mid \in \mathcal{MID}$, $h_1, h_2 \in \mathcal{H}$, $v_k, \ldots, v_{k+n-1} \in \mathcal{V}$, $v_a, \ldots, v_e \in \mathcal{V}$, $as \in \mathcal{Q}$, $u \in \mathcal{V}$, such that $(mid, h(r(v_k))) \in libSpec$ and*

$$\text{ALeakMsg} : \langle \varepsilon, as, \varepsilon \rangle \xrightarrow{(\text{out}, mid, v_a, \ldots, v_e)}_A \langle (mid, v_a, \ldots, v_e), as, \varepsilon \rangle$$

$$\text{AReceiveMsg} : \langle i, as, res \rangle \xrightarrow{(\text{in}, (lo(res), up(res)))}_A \langle \epsilon, as', \epsilon \rangle \qquad \text{for } res \in \mathcal{V}$$

$$\text{AFinal} : \langle i, as, res \rangle \xrightarrow{(\text{final}, res)}_A \langle \epsilon, as', \epsilon \rangle \qquad \text{for } res \in \mathcal{ADVR}$$

for $v_a, v_b, v_c, v_d, v_e \in \mathcal{X}$, $n \in \mathbb{N}$ and $mid \in Mal$,

Figure 11: ADL split-state representation, adversary semantics, ε denotes the empty string.

$m = \mathit{lookup\text{-}direct}_\Pi(mid, h(r(v_k)).\mathsf{class})$,

$$\Pr[\langle (m), h, (0), (\mathit{defReg}([r_k, \ldots, r_{k+n-1}\rangle)), as] \to^* \langle m' \cdot ml, h, pp \cdot ppl, rl, as' \rangle \\ \wedge\, m[pp] = \mathsf{invoke\text{-}static}\ \ v_a, \ldots, v_e, mid \wedge mid \in Mal_{static}] = 0, \text{ and}$$

- *only the Crypto-API semantics and the adversary semantics in* $\text{ADL}^{SS}_{\Pi^\eta, \mathcal{A}^\eta, s_0}$ *make use of the* rand *-instruction, i.e., for any transition* $s_0 \to_{\Pi, n_1, p_1} s_1 \cdots \to_{\Pi, n_l, p_l} s_l$ *to some state* $s_l = \langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as \rangle$ *with* $m[pp] =$ rand v_a *for some* v_a*, there is a position in the method stack, i.e., an* $i \in \mathbb{N}$ *and* $m_i \cdots ml'$*,* $r_i \cdot rl'$ *and* $pp_i \cdot rl'$ *the* i*th suffix of* ml*,* rl *or* ppl*, respectively, such that* $m_i[pp_i] =$ invoke-direct-range v_k, n, mid *and* $(mid, h(r(v_k))) \in dom(libSpec)$*, for some* v_k*,*n *and* mid*, and* h *the heap when this method was called, i.e.,* h *such that, for the largest* $j \in \{0, l\}$ *and some* $as' \in \mathcal{Q}$*,* $s_j = \langle m_i \cdots ml', h, r_i \cdot rl', pp_i \cdot rl', as' \rangle$*.*

Definition 40 (ADL split-state representation)**.** *Given a uniform family of ADL programs* $\Pi = \{\Pi^\eta\}_{\eta \in \mathbb{N}}$*, every member thereof pre-compliant with a crypto-API specification libSpec, a family of attackers* $\{\mathcal{A}^\eta\}_{\eta \in \mathbb{N}} \in \mathcal{ADV}$*, and an initial configuration* $s_0 = \langle ml, h, ppl, rl \rangle$*, we define the ADL split-state representation* $(\text{ADL}^{SS}_{\Pi^\eta, \mathcal{A}^\eta, s_0})_{\eta \in \mathbb{N}}$ *as the family of split-state compositions of the following three transition systems for every* $\eta \in \mathbb{N}$*:*

- honest-program semantics: *The transition system* $(\mathrm{S}_H, s^0_H, A_H, \delta_H)$ *is defined by the ADL semantics(not including the transitions in Figure 7) , extended with rules in Figure 9, and the initial state* $\left\langle ml, h, ppl, rl, as^H_{dmy} \right\rangle$ *for an arbitrary initial adversary state* as^H_{dmy} *which will be ignored.*
- crypto-API semantics: *The transition system* $(\mathrm{S}_L, s^0_L, A_L, \delta_L)$ *is defined by the ADL semantics except* rReturnVF, rReturnF*, extended with rules in Figure 7 and* Prob *(see Figure 10). The initial state is* $\left\langle (), \emptyset, (), (), as^L_{dmy} \right\rangle$ *for some arbitrary initial adversary state* as^L_{dmy} *which will be ignored.*
- attacker semantics: *The transition system* $(\mathrm{S}_A, s^0_A, A_A, \delta_A)$ *consists of the transitions of the adversary* $\mathcal{A}$*, extended with the transitions in Figure 11.*

To make our notation more concise, we abbreviate the initial state $(\top, \left\langle ml, h, ppl, rl, as^H_{dmy} \right\rangle), (\bot, s_A), (\bot, \left\langle (), \emptyset, (), (), as^L_{dmy} \right\rangle)$ *for some honest initial state* $\langle ml, h, ppl, rl \rangle$ *as* $\langle ml, h, ppl, rl \rangle^{ss}$*.*

Observe that the adversarial steps are explicit in this representation, and hence each step is annotated with a computation time of 1.

Definition 41 (Single-step transition system)**.** *A probabilistic transition system is* single-step *if* $s \xrightarrow{[p]}_n s', p > 0$ *or* $s \xrightarrow{a}_n s'$ *imply* $n = 1$*.*

For probabilistic transition steps that are single-step, such as any system that is part of an ADL split-state representation, we will omit the number of computation steps in each transition from here forth.

While the overall split-state representation is probabilistic, the honest program semantics is now deterministic, except for transitions which call the Crypto-API or the adversary. This property is helpful in formulating the embedding, as this means that from any given state, any follow-up state up to the point where the Crypto-API or the adversary is called are uniquely defined.

Definition 42 (Internally deterministic)**.** *A single-step probabilistic transition system $T = (S, s_0, A, \delta)$ is* internally deterministic *if every probabilistic state is deterministic, and if every non-deterministic state s is either deterministic, or if for $A' = \texttt{in}$ and $A' = \texttt{LR}$, there are $\delta_{A'} \colon S \times A' \to S$ such that*

$$\delta_H(s) \subseteq \{(a, s', 1) \mid a \in A' \wedge s' = \delta_{A'}(s, a)\}.$$

Lemma 2. *The honest-program semantics within any ADL split-state representation is internally deterministic.*

Proof. By definition, the honest program semantics exclude probabilistic transitions. By careful inspection, we can verify that all states except those covered by the side-conditions in Figure 9 are deterministic (only one inference rule applies to each state, and each rule determines a unique follow-up state). The rules in Figure 9 are either deterministic, or in case of LibResponse and ReceiveMsg, determine the follow-up state depending on the action. □

Lemma 3. *For all ADL programs Π that are pre-compliant with a crypto-API specification libSpec, all initial configurations $\langle ml, h, ppl, rl\rangle$, all adversaries all $n \in \mathbb{N}$, we have*

$$\Pr[\langle \Pi\, \langle ml, h, ppl, rl\rangle \,\|\mathcal{A}\rangle \downarrow_n x] = \Pr[T \downarrow_n x],$$

where T denotes the ADL split-state representation of Π and $\mathcal{A}$ for initial configuration $\langle ml, h, ppl, rl\rangle$.

Proof. We need to show the equivalence between

$$\sum_{\substack{(s_0 \xrightarrow{\alpha_1}_{n_1} \\ \cdots \xrightarrow{\alpha_k} s_k) \\ \in \\ supp(\mathrm{Exec}(\mathrm{ADL}_{\Pi,\mathcal{A},\langle ml,h,ppl,rl\rangle}))}} \Pr\left[\mathrm{Exec}(\mathrm{ADL}_{\Pi,\mathcal{A},\hat{s_0}}) = s_0 \xrightarrow{\alpha_1}_{n_1} s_1 \cdots \xrightarrow{\alpha_{k-1}}_{n_{k-1}} s_{k-1} \xrightarrow{\alpha_k}_{n_k} \langle x\rangle \wedge \sum_{i=1}^{k} n_i \leq n\right]$$

where $\hat{s_0}$ is the configuration belonging to s_0, i.e., s_0 without the attacker state, and

$$\sum_{(s_0 \xrightarrow{\alpha_1} \cdots s_k) \in supp(\mathrm{Exec}(T))} \Pr\left[\mathrm{Exec}(T) = s_0 \xrightarrow{\alpha_1}_{n_1} s_1 \cdots \xrightarrow{\alpha_k = (\texttt{final}, x)}_{n_k} s_k \wedge k \leq n\right].$$

We will prove this equivalence by induction on the length of the execution in ADL_{Π^n}, however, in order to define what it means for two executions to correspond, we partition execution steps into honest program states, library states, and final states. The initial configuration is an honest program state by definition. Final states stay final states. Each other state is an honest state, if the previous state was an honest state and the transition between the two was not an instance of rIDR where $(mid, h(r(v_k))) \in libSpec$. If the previous state was an honest state, but this condition is not true, the successor state is a library state. Similarly, when the previous state is a library state and the transition is an instance of rReturnV or rReturn such that the state of the successor state matched the stack

of the latest previous state that was not a library state, then the successor is an honest program state, more formally: if $s_i \to s_{i+1}$ is an instance of rReturnV or rReturn, the last previous honest program state s_j with $j < i$ is such that $s_j = \langle ml, h, pp \cdot ppl, rl, as\rangle$, then s_{i+1} is a program state if $s_{i+1} = \langle ml, h', pp + 1 \cdot ppl, rl', as'\rangle$ for some h', rl' and as'.

Given this transition, we define a relation $\mathcal{R}$ between executions, which is, for the most part, based on the resulting state.

- A state s in ADL_{Π^η}, which is a final state, i.e., $s = \langle res\rangle$ for some res, corresponds to a state $s_T = (s_H, s_A, s_L)$ in T^η, if the transition to s_T in T^η was labelled $(\texttt{final}, res)$.
- A state s in ADL_{Π^η}, which is an honest program state of form $s = \langle u, h, as\rangle$, corresponds to a state $s_T = (s_H, s_A, s_L)$ in T^η, if $s_H = (\top, \langle u, h\rangle)$ and $s_A = (\bot, \langle \epsilon, as, \epsilon\rangle)$.
- A state s in ADL_{Π^η}, which is an honest program state $\langle ml, h, ppl, rl, as\rangle$ corresponds to state (s_H, s_A, s_L) iff. $s_H = (\top, \langle ml_H, h_H, ppl_H, rl_H, as_H\rangle)$, $s_L = (\bot, \langle ml_L, h_L, ppl_L, rl_L, as_L\rangle)$, $s_A = (\bot, (in, as, out))$, and
  - $ml = ml_H$,
  - $h = h_H$,
  - $ppl = ppl_H$,
  - $rl = rl_H$.
- A state s in ADL_{Π^η}, which is a library state $\langle ml, h, ppl, rl, as\rangle$ corresponds to state (s_H, s_A, s_L) iff. $s_H = (\bot, \langle ml_H, h_H, ppl_H, rl_H, as_H\rangle)$, $s_L = (\top, \langle ml_L, h_L, ppl_L, rl_L, as_L\rangle)$, $s_A = (\bot, (in, as, out))$, and
  - $ml = ml_L \cdot ml_H$, where $ml_L \neq \epsilon$,
  - $h = h_H[h_L]$,
  - $ppl = ppl_L \cdot ppl_H$,
  - $rl = rl_L \cdot rl_H$.
  - $ml_H = ml_{enter}$, $rl_H = rl_{enter}$ and $ppl = ppl_{enter}$ for $s_i = \langle ml_{enter}, h_i, ppl_{enter}, rl_{enter}, as_i\rangle$, $i < k$ the last program state such that s_{i+1} and all subsequent states were library states

Fix n. We show by induction on number k of steps in execution: For any execution of $k' = \sum_{i=0}^{k} n_i$ steps that ends in a state $\langle ml, h, ppl, rl, as\rangle$ in ADL_{Π^η}, the sum of probabilities of all executions of k' steps in T^η that ends in a corresponding state is exactly the same. Note that for $a \neq a'$, $a\mathcal{R}b$ implies $\neg(a'\mathcal{R}b)$, hence this implies the claim.

Initially, this is the case, because if $\langle ml, h, ppl, rl, (\eta, z)\rangle$ in ADL_{Π^η}, then $s_H = (\top, \langle ml, h, ppl, rl\rangle)$, $s_L = (\bot, (\epsilon, \emptyset, \epsilon, \epsilon, as_{dmy}))$ and $s_A = (\bot, (\eta, z))$. (W.l.o.g., we use as_{dmy} for as^L_{dmy} or as^H_{dmy}, depending on context.)

Assume the induction hypothesis (IH) holds for k steps, i.e., up to a state s_k in ADL_{Π^η}. We perform a case distinction on s_k.

If s_k is a library state, it has the form $s_k = \langle ml, h, ppl, rl, as\rangle$, as the rules rReturnVF and rReturnF are not part of the library semantics. Thus $ml = ml' \cdot ml_{enter}$ for some non-empty sequence ml', as rReturnV, rReturnVF, rReturn and rReturnF are the only rules reducing the method stack, but if one of them would have been applied, s_k would not be a library state. By IH, we have that $ml = ml_L \cdot ml_H$, and $ml_H = ml_{enter}$ so any rule that does not alter the method stack preserves the correspondence, given that Π^η is pre-compliant (see Condition 39). This holds for method calls to honest libraries, too, as they extend ml and, likewise, ml_L to the left. Method call to the adversary are excluded by Condition 39. As s_k's frame is of length at least 2, rReturnVF rReturnF never apply. Consider now the two remaining cases, where either rReturnV or rReturn are applied. We treat

only the latter, as the former is completely analogous. Let $s_k = \langle m \cdot ml', h, pp \cdot ppl, r \cdot r' \cdot rl, as\rangle$ for $ml' \neq \epsilon$ and $m[pp] = \mathsf{return}\ v_a$. As $m \cdot ml' = ml_L \cdot ml_{enter}$, either $ml_L = (m)$ and thus $ml' = ml_{enter}$, or $|ml_L| > 1$. In the first case, $ml_L = (m)$, only LLibRetVoid or LLibRet (depending in $m[pp]$) apply. These are labelled $(\texttt{LR}, lo(r(v_a)), up(r(v_a)), h)$ with probability 1 and 1 computation step. By asynchronous parallel composition, the the honest program semantics moves with rule LibResponse. The following state is thus an honest program state, and relation holds. In particular, the library's heap overwrites the honest program's, thus it holds that $h' = h$ (by definition of the transition), $h = h_H[h_L]$ (by IH) $h_H[h_L] = h'_H$ (by LibResponse) and thus $h' = h'_H$. In the second case, $|ml_L| > 1$, rReturnV (or rReturn) are used in both cases, preserving the IH. It is not possible to enter a final state from s_k, as rReturnVF and rReturnF only apply if the method stack is of size 1.

If s_k is an honest program state, it has the form $s_k = \langle ml, h, ppl, rl, as\rangle$, or $s_k = \langle u, h, as\rangle$. In the second case, $s_k = \langle u, h, as\rangle$, the only applicable rule is AdvFin, resulting in a transition labelled $(\texttt{final}, res)$. to state $\langle res\rangle$, for the sum of all probabilities of adversary steps that go from $\mathcal{A}$ to res in n steps, for some n. Let p be this sum. The only applicable rule in T^η is FinalCall, which has label $(\text{out}, \texttt{finalCall})$ and thus executes ALeakMsg in parallel. There is a set of corresponding executions of the adversary for each n_k, the sum of which has probability p, as the number of steps, starting states and final step are the same. Including the final step AFinal (which is always applicable when the third element of the triple is in $\mathcal{ADVR}$, and excludes AReceiveMsg), each of these results in a transition $(\texttt{final}, res)$. for $n+2$ steps. Now we consider the first case, $s_k = \langle ml, h, ppl, rl, as\rangle$. Let the transition to s_{k+1} be an instance of rIDR where $(mid, h(r(v_k))) \in libSpec$. Then, the next state is a library state $sk' = \langle m' \cdot ml, h, 0 \cdot ppl, defReg([r(v_k), \ldots, r(v_{k+n-1})\rangle) \cdot rl, as]$ For each corresponding state in T^η, LibCall in parallel with LLibCall is the only transition possible in T^η. By asynchronous parallel composition, each follow-up state will be such that $s_H = (\bot, ml, h, ppl, rl, as')$, $s_L = (\top, (m), h|_{r(v_k)}, (0), defReg([r(v_k), \ldots, r(v_{k+n-1})]), as'')$, and $s_A = (\bot, in, as, out$ for some $as', as'' \in \mathcal{ADVR}$, $in \in (F_{mal}, \mathcal{V}^*)$, and $out \in \times(\mathcal{V} \cup \mathcal{ADVR} \cup \{\epsilon\})$. These library states correspond to each other, in particular, $h = h_H = h_H[h_L] = h_H[h|_{r(v_a)}]$, and $ml = ml_{enter}$, $ppl = ppl_{enter}$, $rl = rl_{enter}$. Let the transition be an instance of the rule AdvInv with probability p in n steps. Similar to the previous case involving the attacker, the sum of all possible transitions in $n+1$ steps from (s_H, s_A, s_L), equals p, as only rules LeakMsg and ALeakMsg apply (in parallel). The adversary is activated on same input, i.e., all instances of states $(i, as, res$ with $res \in \mathcal{V}$, reachable in T synchronise with ReceiveMsg by definition of AReceiveMsg and Definition 22. The number of computation steps account for the extra activation with ReceiveMsg in parallel with ReceiveMsg. Let the transition be an instance of the rule AdvInv with probability p in n steps. Similar to the previous case, where AdvInv was invoked from a library state, this transition is labelled $(\texttt{final}, res)$. and results in a state $\langle res\rangle$, for the sum of all probabilities of adversary steps that go from $\mathcal{A}$ to res in n steps, for some n. The argument remains the same. Note that the adversary produces a final state (which cannot have successor, by Definition 22), thus there is no final state before any final state. Hence LeakMsg and ALeakMsg synchronise as before, and AFinal is the only way to proceed from final state. If the transition was an instance of Prob, by Condition 39, Definition 39, s_k would not be an honest-program state. Any other transition affects only the left-most element of ml, ppl and rl, thus the same transition can be applies in T^η.

If s_k is a final state, there is no next step in s_k and all corresponding states. This concludes the proof. □

7.3 Over-approximating ADL

Recall that our plan is to instantiate $\mathcal{D}_c$ with symbolic terms and even with positions in a CoSP tree. However, some operations are not defined on symbolic terms. As an example consider the XOR

operation. It has been shown that each computationally sound symbolic representation of XOR has to work on symbolic bitstrings. Hence, for a typical symbolic representation of a ciphertext the XOR operation is not defined in the symbolic model.

In order to get rid of any undefined operations after instantiating $\mathcal{D}_c$ with symbolic terms or CoSP tree positions, we over-approximate each undefined operation by querying the attacker for the result. To this end, we define an over-approximation of the ADL semantics in terms of its split-state representation. This over-approximation tracks values resulting from calls to the crypto-API. Whenever a computation, e.g., a unary operation, a binary operation or a test, is performed on these values, the over-approximation gives the adversary more power: she can decide the values of these computations. This is necessary, as these operations cannot be computed in the symbolic model. Already performing this over-approximation on the ADL split-state semantics simplifies the embedding. Thanks to the split-state composition, we can define this over-approximation canonically.

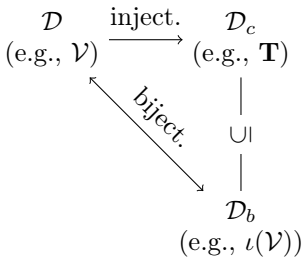

Figure 12: Domains of the canonical over-approximation.

We assume some domain $\mathcal{D}$ underlying the transition function and the set of states of the honest program semantic, and a distinct new domain $\mathcal{D}_c$ (for values resulting from the crypto API), a subset of which $\mathcal{D}_b \subset \mathcal{D}_c$ (representation of bitstrings) have a bijection to the original domain (see Figure 12). We assume the honest-program semantics to be defined via a set of rules, and interpret these rules in the new domain $\mathcal{D}_c$. If the re-interpreted rule can be instantiated regardless of whether some state carries values in $\mathcal{D}$ or $\mathcal{D}_c$, this rule is transferred to the over-approximated semantics. Most rules only move values from registers to heaps and are homomorphic in this sense. Rules for which the above do not hold, but which can be expressed using a constructor or destructor, are split into four rules, two of which send these values to the adversary and use her input for the follow-up state in case one of the variables is in $\mathcal{D}_c \setminus \mathcal{D}_b$. Otherwise, i.e., if only values from $\mathcal{D}$ are used, or values representable in $\mathcal{D}$, then the crypto-API is used for the computation. This way, we were able to encode binary operations like XOR as destructors, so representations of bitstrings that have passed the crypto API (e.g., a bitstring was encrypted and then decrypted again) can be treated without unnecessary imprecision. If any rule $r \in R_H$ falls in neither of the above cases, the canonical over-approximation is undefined.

Using this over-approximation, we can derive ADLo^{SS}, the over-approximated ADL semantics. We tag messages resulting from calls to the crypto-API using a set $\mathcal{N}_c$, such that $\mathcal{N}_c \cap \mathcal{N} = \emptyset$, and assume a bijection between the two that is efficiently computable. For $n \in \mathcal{N}_c \cup \mathcal{N}$, let $[n]_{\mathcal{N}_c}$ and $[n]_{\mathcal{N}}$ denote its representation in $\mathcal{N}_c$ or $\mathcal{N}$ according to the bijection. We lift this notation to values in $\mathcal{V}$, too. The inference rules include those previously defined, but modified such that register values are converted to $\mathcal{N}$ before addressing the crypto-API, i.e., $\texttt{LC}(f, r(v_a), \ldots, r(v_e))$ is substituted by $\texttt{LC}(f, [r(v_a)]_{\mathcal{N}_c}, \ldots, [r(v_e)]_{\mathcal{N}_c})$, and $\texttt{out}(mid, r(v_a), \ldots, r(v_e))$ is substituted by $\texttt{out}(mid, [r(v_a)]_{\mathcal{N}_c}, \ldots, [r(v_e)]_{\mathcal{N}_c})$. Finally, the rules in Figure 13 are added.

Precise treatment of bitstring representations output by Crypto-API. If some value was produced by the Crypto-API, e.g., an encryption, in ADLs, it will be represented by a symbolic term, e.g., $enc(k, m)$ for some terms k and m. When $enc(k, m)$ is tested for equality to a bitstring, there is no meaningful way to decide for the outcome of this test. To simplify the next proof steps, we

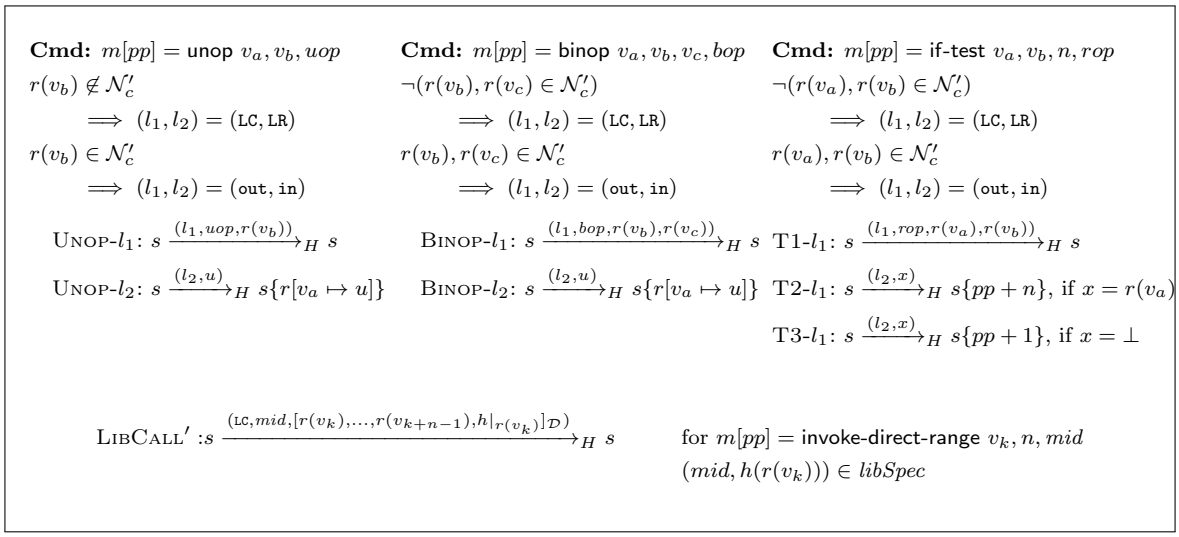
Cmd: $m[pp] = \mathsf{unop}\ v_a, v_b, uop$
$r(v_b) \notin \mathcal{N}'_c$
$\implies (l_1, l_2) = (\mathtt{LC}, \mathtt{LR})$
$r(v_b) \in \mathcal{N}'_c$
$\implies (l_1, l_2) = (\mathtt{out}, \mathtt{in})$
Unop-l_1: $s \xrightarrow{(l_1, uop, r(v_b))}_H s$
Unop-l_2: $s \xrightarrow{(l_2, u)}_H s\{r[v_a \mapsto u]\}$
Cmd: $m[pp] = \mathsf{binop}\ v_a, v_b, v_c, bop$
$\neg(r(v_b), r(v_c) \in \mathcal{N}'_c)$
$\implies (l_1, l_2) = (\mathtt{LC}, \mathtt{LR})$
$r(v_b), r(v_c) \in \mathcal{N}'_c$
$\implies (l_1, l_2) = (\mathtt{out}, \mathtt{in})$
Binop-l_1: $s \xrightarrow{(l_1, bop, r(v_b), r(v_c))}_H s$
Binop-l_2: $s \xrightarrow{(l_2, u)}_H s\{r[v_a \mapsto u]\}$
Cmd: $m[pp] = \mathsf{if\text{-}test}\ v_a, v_b, n, rop$
$\neg(r(v_a), r(v_b) \in \mathcal{N}'_c)$
$\implies (l_1, l_2) = (\mathtt{LC}, \mathtt{LR})$
$r(v_a), r(v_b) \in \mathcal{N}'_c$
$\implies (l_1, l_2) = (\mathtt{out}, \mathtt{in})$
T1-l_1: $s \xrightarrow{(l_1, rop, r(v_a), r(v_b))}_H s$
T2-l_1: $s \xrightarrow{(l_2, x)}_H s\{pp + n\}$, if $x = r(v_a)$
T3-l_1: $s \xrightarrow{(l_2, x)}_H s\{pp + 1\}$, if $x = \bot$
LibCall′ :$s \xrightarrow{(\mathtt{LC}, mid, [r(v_k), \ldots, r(v_{k+n-1}), h|_{r(v_k)}]_{\mathcal{D}})}_H s$
for $m[pp] = \mathsf{invoke\text{-}direct\text{-}range}\ v_k, n, mid$
$(mid, h(r(v_k))) \in libSpec$

Figure 13: over-approximated honest program semantics for ADL (ADLo^{SS}). Here $\mathcal{N}_b := range(\iota)$ and $\mathcal{N}'_c := \mathcal{N}_c \setminus \mathcal{N}_b$.

over-approximate in these situations, even though we have not made the transition into the symbolic model yet. But not every output from the Crypto-API should be over-approximated in this. Consider the following example: a bitstring is encrypted by the Crypto-API, stored and later decrypted again. The result of this decryption is very well a library output, but it is perfectly feasible to compute equivalence to a second bitstring. It should not be necessary to send the bitstring to the adversary in this case. In terms of the symbolic model, the bitstring would be represented as outlined in Example 2 on page 25. Each bitstring has a unique term representation, thus there exists a subset of $\mathbf{T}$ which is bijective to the set of bitstrings. In general, we assume a subset of $\mathcal{D}_c$, called $\mathcal{D}_b$, which is injective to the original domain $\mathcal{D}$. If an operation, e.g., a binary operation like XOR, is to be applied to two values in $\mathcal{D}_b$, we require the library to compute this operation. For all relevant operations in ADL, this is possible w.l.o.g (see Appendix A. We thus obtain a relatively precise over-approximation.

Rules. We assume the transition relation to be defined in terms of a set of inference rules employing (meta-language) variables parametric in some domain $\mathcal{D}$. We say we interpret a rule in the domain $\mathcal{D}'$, if we alter the domain of each variable in $\mathcal{D}$ to $\mathcal{D}'$. Naturally, this might lead to rules that cannot be satisfied any more, e.g., if $\mathcal{D} = \mathbb{R}$ and a premise requires a variable $v \in \mathbb{R}$ to equal π, the same rules interpreted in the domain $\mathcal{D}' = \mathbb{N}$ are unsatisfiable. However, most rules defining the ADL semantics can be re-interpreted easily. Many of them move values from one register to another, or from heap to registers and are thus oblivious of the actual type of data moved. If a transition from a state s to a state s' is possible with, e.g., rMove, and there is a well-defined mapping from states in the domain $\mathcal{D} = \mathcal{V}$ to states in any other domain $\mathcal{D}'$, then an interpretation of rMove in $\mathcal{D}$' can be instantiated for s and s' mapped to $\mathcal{D}'$.

Some ADL rules, e.g., rBinop, cannot be re-interpreted for an arbitrary $\mathcal{D}$. In this case, the operation $\underline{bop}$ is only defined on values in $\mathcal{V}$, more precisely $\mathcal{N}$. Hence the predicate $x = r(v_a)\ \underline{bop}\ r(v_b)$ is unsatisfiable if $r(v_a), r(v_b) \notin \mathcal{V}$. As mentioned before, we can deal with this, by *a*) transforming $r(v_a)$ and $r(v_b)$ to $\mathcal{D}_b$ and sending them to the library, if they are in $\mathcal{D}$ or $\mathcal{D}_b$, or *b*) sending them $r(v_a)$ and $r(v_b)$ if they are in $\mathcal{D}_c \setminus \mathcal{D}_b$, i.e., if it is not clear how to evaluate this operation.

We define this over-approximation for any semantics in split-state form as follows. Here we benefit from the split-state composition, as it allows for rewriting rules in a way that involves the attacker, or the Crypto-API. The over-approximated ADL semantics is an instance of the canonical over-

approximation.

Definition 43 (Crypto-API space)**.** *Let $\mathcal{D}$ be any set. A pair of domains $(\mathcal{D}_c, \mathcal{D}_b)$ is a* Crypto-API space*, if*

- $\mathcal{D}_c \cap \mathcal{D} = \emptyset$,
- *there is an injection f from $\mathcal{D}$ to $\mathcal{D}_c$, and*
- $\mathcal{D}_b \subset \mathcal{D}_c$, *and*
- *there is a bijection b between $\mathcal{D}$ and $\mathcal{D}_b$.*

In addition, we use the following notation for all $d \in \mathcal{D} \cup \mathcal{D}_b$:

$$[d]_{\mathcal{D}} = \begin{cases} d & \text{if } d \in \mathcal{D} \\ b^{-1}(d) & \text{if } d \in \mathcal{D}_b \end{cases} \qquad [d]_{\mathcal{D}_b} = \begin{cases} b(d) & \text{if } d \in \mathcal{D} \\ d & \text{if } d \in \mathcal{D}_b \end{cases} \qquad [d]_{\mathcal{D}_c} = \begin{cases} f(d) & \text{if } d \in \mathcal{D} \\ d & \text{if } d \in \mathcal{D}_c \end{cases}$$

We lift this notation to finite sets, tuples and sequences.

Definition 44 (Canonical over-approximation (honest program sementics))**.** *Let $(\mathcal{D}_c, \mathcal{D}_b)$ be a Crypto-API space for $\mathcal{D}$ and $\mathcal{D}' := \mathcal{D} \cup \mathcal{D}_b \cup \mathcal{D}_c$. Given a symbolic model $\mathbf{M} = (\mathbf{C}, \mathbf{N}, \mathbf{T}, \mathbf{D})$, and an internally deterministic transition system T_H with $\rightarrow_H$ defined in terms of a finite set of inference rules R_H, which are parametric in a domain $\mathcal{D}$, we define $\rightarrow_{H'}$ from the set of inference rules R_H that defines $\rightarrow_H$:*

(i) Each rule $r \in R_H$, for which any instance with domain $\mathcal{D}$ that concludes $s \xrightarrow{(l,1,1)}_H s'$ for some s and s', can be instantiated with domain $\mathcal{D}_b$ to conclude $[s]_{\mathcal{D}_b} \xrightarrow{(l,1,1)}_{H'} [s']_{\mathcal{D}_b}$, and with domain $\mathcal{D}_c$ to conclude $[s]_{\mathcal{D}_c} \xrightarrow{(l,1,1)}_{H'} [s']_{\mathcal{D}_c}$, is part of the rules when interpreted with domain $\mathcal{D}_c$ (which $\mathcal{D}_b$ is a subset of).

(ii) Each rule $R \in R_H$, for which the above does not hold, but for which there is a symbolic operation $O \in \mathrm{SO}$ such that it has the form

$$R = \frac{\phi(s) \wedge \rho(s, s') \wedge \Pr\left[[r]_{\mathcal{D}_b} = A_O([v_1, \ldots, v_l]_{\mathcal{D}_b})\right] = 1}{s \xrightarrow{(\varepsilon,1,1)}_H s'}$$

for variables $v_1, \ldots, v_l$ in s, and variables r in s', we define four rules:

$$\frac{\phi([s]_{\mathcal{D}}) \wedge \neg(v_1, \ldots, v_l \in \mathcal{D} \cup \mathcal{D}_b)}{s \xrightarrow{((\mathtt{out}, O, v_1, \ldots, v_l),1,1)}_{H'} s} \qquad \frac{\phi([s]_{\mathcal{D}}) \wedge \rho'([s]_{\mathcal{D}}, [s']_{\mathcal{D}}) \wedge \neg(v_1, \ldots, v_l \in \mathcal{D} \cup \mathcal{D}_b)}{s \xrightarrow{((\mathtt{in}, r),1,1)}_{H'} \delta_{\mathtt{in}}(s, r) = s'}$$

$$\frac{\phi([s]_{\mathcal{D}}) \wedge v_1, \ldots, v_l \in \mathcal{D} \cup \mathcal{D}_b}{s \xrightarrow{((\mathtt{LC}, O, [v_1]_{\mathcal{D}_b}, \ldots, [v_l]_{\mathcal{D}_b}),1,1)}_{H'} s} \qquad \frac{\phi([s]_{\mathcal{D}}) \wedge \rho([s]_{\mathcal{D}}, [s']_{\mathcal{D}}) \wedge v_1, \ldots, v_l \in \mathcal{D} \cup \mathcal{D}_b}{s \xrightarrow{((\mathtt{LR}, r),1,1)}_{H'} \delta_{\mathtt{LR}}(s, r) = s'}$$

If any rule $r \in R_H$ falls in neither of the above cases, the canonical over-approximation is undefined. Otherwise, the canonical over-approximation of the honest program semantics T_H for Crypto-API space $(\mathcal{D}_c, \mathcal{D}_b)$ *is the transition system*

$$T'_H := ([\mathrm{S}_H]_{\mathcal{D}'}, [\mathrm{S}^0_H]_{\mathcal{D}'}, \rightarrow_{H'}).$$

Definition 45 (Canonical over-approximation)**.** *Given a symbolic model* $\mathbf{M} = (\mathbf{C}, \mathbf{N}, \mathbf{T}, \mathbf{D})$*, and a split-state composition* T *of an internally deterministic transition system* T_H*, and transition systems* T_A *and* T_L*, such that* T *is probabilistic, let* T'_H*, the canonical over-approximation of the honest program sementics* T_H*, be defined.*

Then, the canonical over-approximation of T for the Crypto-API space $(\mathcal{D}_c, \mathcal{D}_b)$ *is the split-state composition of* T'_H*,* T_A*,* T'_L*, where*

$$T'_H := ([\mathrm{S}_H]_{\mathcal{D}'}, [\mathrm{S}^0_H]_{\mathcal{D}'}, \rightarrow_{H'})$$
$$T'_L := (\mathrm{S}_L, \mathrm{S}^0_L, \rightarrow_{L'})$$

with $s \xrightarrow{(\text{LC}, f, [v_1]_{\mathcal{D}_b}, \ldots, [v_l]_{\mathcal{D}_b})}_{L'} s'$ *if* $s \xrightarrow{(\text{LC}, f, v_1, \ldots, v_l)}_{L} s'$ *and* $s \xrightarrow{\alpha}_{L'} s'$ *if* $s \xrightarrow{\alpha}_{L} s'$ *and* $\alpha \notin$ libCall.

Lemma 4. *For any split-state composition* T *from an honest program* T_H*, and some attacker and library transition system, the honest program semantics* $T'_H = (\mathrm{S}'_H, \mathrm{S}^{0'}_H, \rightarrow'_H)$ *resulting from the canonical over-approximation of* T *to the crypto-API space* $(\mathcal{D}_c, \mathcal{D}_b)$ *are internally deterministic, if* T_H *is internally deterministic.*

Proof. Consider the subset of the rules defining $\rightarrow'_H$ that are just re-interpreted according to Condition i, Definition 45. If the transition relation relating from this subset was not internally deterministic, T_H would not have been internally deterministic, as the mapping from $\mathcal{D}$ to $\mathcal{D}_c$ is injective. Consider now any additional rule derived according to Condition ii, Definition 45. Either the first and second rule applies, or the third and fourth, as their premises are mutually exclusive on the state. By definition of these rules, the follow-up state s' depends only on s and value $r \in \mathcal{D}'$ contained in the label. It is left to show that no state s to which any of these four rules can be instantiated, is covered by any rule resulting included according to Condition i, Definition 45, or by another rule resulting fro Condition ii. In the both cases, this would imply that T_H was not internally deterministic, in the first case directly, in the second case because the same predicate on s is in the premise of all four translated rules. □

Lemma 5. *For any split-state composition* T *from an honest program* $T_H = (\mathrm{S}_H, \mathrm{S}^0_H, \rightarrow_H)$*, an attacker* T_A*, and a library transition system* $T_L = (\mathrm{S}_L, \mathrm{S}^0_L, \rightarrow_L)$*, if*

- $\Pr[s \xrightarrow{(\text{out}, O, u_1, \ldots, u_l)}_H s'] = 0$ *and* $\Pr[s \xrightarrow{(\text{LC}, O, u_1, \ldots, u_l)}_H s'] = 0$ *for any two states* $s, s' \in \mathrm{S}_H$ *and any sequence* $u_1, \ldots, u_l$*, and*
- $\Pr[s \xrightarrow{(\text{LC}, O, v_1, \ldots, v_l)}_L \rightarrow^* \xrightarrow{(\text{LR}, [A_O(v_1, \ldots, v_l)]_{\mathcal{D}})}_L s'] = 1$ *for any* $s' \in \mathrm{S}_L$ *and any* s *for which holds* $s \in \mathrm{S}^0_L$ *or* $\Pr[s' \xrightarrow{\text{libResp}} s] \neq 0$*.*

then there is an attacker T'_A *such that for the canonical over-approximation* T' *from* T_H, T_L *and* T'_A*, and for all* $\eta, n \in \mathbb{N}$*,* $z \in \{0,1\}^*$*,*

$$\Pr[T(\langle ml, h, ppl, rl\rangle, (\bot, (\eta, z)), \langle (), \emptyset, (), (), as_{dmy}\rangle) \downarrow_n x] = \Pr[T'([\langle ml, h, ppl, rl\rangle]_{\mathcal{D}_b}, (\bot, (\eta, z)), ((), \emptyset, ()(), as_{dmy})) \downarrow_{p(n)} x].$$

for some polynomial p*.*

Proof. Let T'_A be T_A, with all transitions of form $(\text{out}, O, \ldots)$ for $O \in$ SO removed, and all transitions labelled $(\text{libCall}, f, u_1, \ldots, u_l)$ for $f \notin$ SO interpreted with $[u_1, \ldots, u_l]_{\mathcal{D}}$ instead of $u_1, \ldots, u_l$ in the successor state. We show that for any transition from step (s_H, s_A, s_L) to $(s'_H, s'_A, s'_L$ in T, there is a sequence of transitions with the same overall probability (i.e., the product of their probabilities) from $([s_H]_{\mathcal{D}_b}, s_A, s_L)$ to $([s']_{\mathcal{D}_b}, s'_A, s'_L)$ in T'. Initially, this holds by definition of the initial state. Inductive

step. Let the honest program be active. For any transition described by the transformation i, this holds immediately, unless the adversary or the library is called. If the adversary is called, observe that, by assumption one, the adversary is never invoked on $(\text{out}, O, \ldots)$ for $O \in \text{SO}$, so T_A differs only in the transformation, from $\mathcal{D}_b$ to $\mathcal{D}$, which, by bijection, produces the same follow-up state. If the library is called, the same argument is made by definition of the library semantics of the over-approximation from Definition 45. transformation ii, consider the instantiation of R that applies from s_H to s'_H. By definition of $[\cdot]_{\mathcal{D}_b}$, the first two rules derived from R never apply, but the third applies. By assumption two and the definition of split-state composition, T' proceeds with probability 1 until the fourth rule applies. In particular, the same predicates hold, and, as in the previous case, by the fact that $\mathcal{D}$ and $\mathcal{D}_b$ are bijective, and by definition of the over-approximated library semantics in Definition 45, the library transitions operate on the same internal states. The resulting state contains $r \in \mathcal{D}$, i.e., the original domain, which is left intact by the conversation $[s']_{\mathcal{D}}$. In the last case, the library or the attacker is active. If no message is passed to the honest program, the the induction hypothesis is trivially preserved. If a message is passed to the honest program, the honest program needs to have a transition with a label in in or libResp, and as T_H needs to be internally deterministic for the over-approximation to be defined, this transition is an instance of a rule that matches Condition i. By the transformation applied to these rules, the values in $\mathcal{D}$ are correctly translated to on reception $\mathcal{D}_b$. □

Definition 46 (Harmonizing Crypto-API)**.** *A Crypto-API* $T_L := (\text{S}_L, \text{S}_L^0, \to_L)$ harmonizes *with a computational implementation* Impl *of a symbolic model* $\mathbf{M}$ *and a libSpec, if*

$$\Pr\left[\exists s' \in \text{S}_L^0.s \xrightarrow{((\text{LC}, op, u_1, \ldots, u_l), 1, 1)} L' \to^* \xrightarrow{((\text{LR}, r), 1, 1)}_L s'\right] = \Pr\left[A_{libSpec(op)}(u_1, \ldots, u_l) = r\right]$$

and

$$\Pr\left[\exists s' \in \text{S}_L^0.s \xrightarrow{((\text{LC}, f, u_1, \ldots, u_l, h), 1, 1)} L' \to^* \xrightarrow{((\text{LR}, r), 1, 1)}_L s'\right] = \Pr\left[A_{libSpec(f, h(u_1))}(u_1, \ldots, u_l, h) = r\right]$$

for any s *for which* $s \in \text{S}_L^0$ *or* $\Pr[s' \xrightarrow{\text{libResp}} s] \neq 0$.

Definition 47 (Crypto-API compliance)**.** *An ADL program* Π *is compliant with a library specification libSpec w.r.t. a symbolic model* $\mathbf{M}$ *and a computational implementation of* $\mathbf{M}$ *called* Impl*, if* Π *is pre-compliant (see Definition 39), and there is a Crypto-API* $T_L := (\text{S}_L, \text{S}_L^0, \to_L)$ *such that* T_L *harmonizes with* Impl *(see Definition 46).*

Example 3: ADLo: over-approximated ADL semantics. We use the canonical over approximation to define ADLo^{SS}, the over-approximated version of the ADL split-state semantics. Let $\mathcal{D} = \mathcal{V} \cup \mathcal{H}$, and $(\mathcal{D}_c, \mathcal{D}_b)$ any Crypto-API space for $\mathcal{D}$. This can, for example, created by tagging library output, and thus having $\mathcal{D}_b = \mathcal{D}_c$. Given a uniform family of ADL programs $\Pi = \{\Pi^\eta\}_{\eta \in \mathbb{N}}$, every member of which is pre-compliant with a library specification *libSpec*, and a family of attackers $\mathcal{A} = \{\mathcal{A}^\eta\}_{\eta \in \mathbb{N}} \in \mathcal{ADV}$, we call the canonical over-approximation of the ADL split-state representation (see Definition 40) the over-approximated ADL semantics, denoted $\text{ADLo}^{SS}_{\Pi, \mathcal{A}, s_0} = (\text{ADLo}^{SS}_{\Pi^\eta, \mathcal{A}^\eta, s_0})_{\eta \in \mathbb{N}}$. ◇

Corollary 1. *For any ADL program* Π *compliant with a library specification libSpec and a symbolic model* $\mathbf{M}$*, and an attacker* $\mathcal{A}$*, there is an attacker* $\mathcal{A}'$ *and a polynomial* p *such that for all* n*:*

$$\Pr[\text{ADL}^{SS}_{\Pi, \mathcal{A}, s_0} \downarrow_n x] = \Pr[\text{ADLo}^{SS}_{\Pi, \mathcal{A}', s_0} \downarrow_{p(n)} x].$$

We define the the split-state equivalence for ADL.

Definition 48 (ADL split-state equivalence)**.** *Let* Π_1 *and* Π_2 *be two families of ADL programs compliant with the same library specification libSpec and* s_1, s_2 *initial configurations for* Π_1 *and* Π_2*, respectively. We write* $\Pi_1\langle s_1\rangle \approx_c^{\text{SS}} \Pi_2\langle s_2\rangle$*, if, for all adversaries* $\mathcal{A} \in \mathcal{ADV}$*, for the over-approximated ADL split-state representations* $(\text{ADLo}^{SS}_{\Pi_1^\eta, \mathcal{A}^\eta, s_1})_{\eta \in \mathbb{N}}$ *of* $\Pi_1\langle s_1\rangle$ *and* $\mathcal{A}$*, and* $(\text{ADLo}^{SS}_{\Pi_2^\eta, \mathcal{A}^\eta, s_2})_{\eta \in \mathbb{N}}$ *of* $\Pi_2\langle s_2\rangle$ *and* $\mathcal{A}$*, we have* $T_{H,1}(s_1) \approx_c^{T_L, T_A} T_{H,2}(s_2)$*.*

Finally, we are in a position to connect the ADL semantics to the over-approximated split-state form.

Lemma 6. *Let Π_1 and Π_2 be two families of ADL programs pre-compliant with the same library specification libSpec and s_1, s_2 initial configuration for Π_1 and Π_2, respectively. Then $\Pi_1\langle s_1\rangle \approx_c^{\mathrm{SS}} \Pi_2\langle s_2\rangle \implies \Pi_1\langle s_1\rangle \approx_c^{\mathrm{ADL}} \Pi_2\langle s_2\rangle$.*

Proof. Follows from Lemma 3 and Corollary 1. □

7.4 Canonical symbolic semantics

The canonical over-approximated split-state semantics define a canonical symbolic semantics, when the Crypto-API space $\mathcal{D}_c$ is instantiated with the set of terms defined by the CoSP-symbolic model $\mathbf{M}$ (in which our result in parametric), when the attacker semantics is given by the symbolic attacker from $\mathbf{M}$, and when the crypto-API is simplified to the constructors and destructors in $\mathbf{M}$.

Definition 49 (Canonical symbolic split-state semantics)**.** *Given a symbolic model $\mathbf{M} = (\mathbf{C}, \mathbf{N}, \mathbf{T}, \mathbf{D})$, a subset $\mathbf{T}_b \subset T$ such that $(\mathbf{T}, \mathbf{T}_b)$ constitute a Crypto-API space, a split-state semantics T with an honest-program semantics T_H, and an attacker strategy, i.e., a sequence $I \in AttS$ for $AttS = \{(\mathtt{in}, O) \mid O \in \mathrm{SO}\}$ (see Definition 6), the* canonical symbolic split-state semantics of T w.r.t. $(\mathbf{T}, \mathbf{T}_b)$ and I *is the split-state composition of*

- *the canonical over-approximation of the honest program semantics for Crypto-API space $(\mathbf{T}, \mathbf{T}_b)$ (see Definition 44),*
- *the attacker semantics $(\mathrm{S}_I, \mathrm{S}_I^0, \to_I)$ with state space $\mathrm{S}_I := AttS \times Event_{out}^* \times (\mathbf{T} \cup \{\bot\})$, initial state $\mathrm{S}_I^0 = (I, \epsilon, \bot)$ ($\epsilon \in Event_{out}^*$), and $\to_I$ the smallest relation such that*

$$((\mathtt{in}, O) \cdot IL, V, \bot) \xrightarrow{(\mathtt{out}, mid, u_1, \ldots, u_l)}_{I,1,1} (IL, V', eval_O(Out(V'))) \quad \text{for } V' = V \cdot (\mathtt{out}, \widehat{mid}, \widehat{u_1}, \ldots, \widehat{u_l})$$

$$(IL, V, pair(u_{lo}, u_{up})) \xrightarrow{(\mathtt{in}, u_{lo}, u_{up})}_{I,1,1} (IL, V, \bot),$$

- *and the Crypto-API semantics $(\mathrm{S}_L, \mathrm{S}_L^0, \to_L)$ with state space $\mathrm{S}_L = \mathbf{T}^3 \cup \{\bot, \epsilon\}$, initial state space $\mathrm{S}_L^0 = \epsilon$ and $\to_L$ the smallest relation such that:*

$$\epsilon \xrightarrow{(\mathtt{LC}, mid, u_1, \ldots, u_l, h)}_{L,1,1} (t_{lo}, t_{up}, t_h,) \qquad \text{if } pair(t_{lo}, pair(t_{up}, t_h)) = eval_O(\widehat{mid}, \widehat{u_1}, \ldots, \widehat{u_l}, \widehat{h})$$

$$\text{or } t_{lo} = t_{up} = t_h = \bot \text{ if } eval_O(\widehat{mid}, \widehat{u_1}, \ldots, \widehat{u_l}, \widehat{h}) = \bot$$

$$(t_{lo}, t_{up}, t_h) \xrightarrow{(\mathtt{LR}, (t_{lo}, t_{up}, t_h))}_{L,1,1} \epsilon.$$

We can define the ADL symbolic split-state semantics as an instance of the canonical symbolic semantics.

Definition 50 (ADLs^{SS}, ADL symbolic split-state semantics)**.** *Let $\mathbf{T}_b$ be the domain of ι (see Definition 29) and observe that $(\mathbf{T}, \mathbf{T}_b)$ constitute a Crypto-API space. For Π a uniform family of ADL programs, every member of which is pre-compliant with a library specification libSpec, the* ADL symbolic split-state semantics of Π w.r.t. an attacker strategy I *is the canonical symbolic split-state semantics of the ADL split-state representation (see 40) w.r.t. $(\mathbf{T}, \mathbf{T}_b)$ and I.*

Note that, given a program Π and an attacker strategy I, the ADL symbolic split-state semantics of Π w.r.t. I is deterministic.

We show that the ADL symbolic split-state semantics can be trivially simplified to the semantics presented in Section 6.

We define symbolic equivalence in the spirit of Definition 9 with equivalence of views defined as in Definition 8.

Definition 51 (Symbolic equivalence $\approx_s^{\mathrm{SS}}$)**.** *Two ADL programs* Π_1 *and* Π_2*, and initial configurations* $s_1 = \langle ml_1, h_1, ppl_1, rl_1\rangle$ *and* $s_2 = \langle ml_2, h_2, ppl_2, rl_2\rangle$ *are* symbolically split-state equivalent ($\Pi_1\langle s_1\rangle \approx_s^{\mathrm{SS}} \Pi_2\langle s_2\rangle$) *if for all attacker strategies* I*, their respective ADL symbolic split-state semantics* T_1, T_2 *w.r.t. to* I *are symbolically equivalent, i.e., if* $\mathrm{SViews}(T_1) \sim \mathrm{SViews}(T_2)$.

Now we are ready to state that symbolic equivalence in the split-state setting implies symbolic equivalence in the sense of Definition 9, or in other words, symbolic equivalence with respect to the semantics introduced in Section 6.

Lemma 7. *Let* Π_1 *and* Π_2 *be two ADL program pre-compliant to the same library specification libSpec and* s_1, s_2 *initial configuration for* Π_1 *and* Π_2*, respectively. Then,*

$$\Pi_1\langle s_1\rangle \approx_s^{\mathrm{ADL}} \Pi_2\langle s_2\rangle \implies \Pi_1\langle s_1\rangle \approx_s^{\mathrm{SS}} \Pi_2\langle s_2\rangle$$

Proof. Assume $\Pi_1 \not\approx_s^{\mathrm{SS}} \Pi_2$, then there are initial configuration $s_1 \in T_{\Pi_1}$, $s_2 \in T_{\Pi_2}$ and views $t_1 \in \mathrm{SViews}_{s_1}(T_{\Pi_1})$, $t_2 \in \mathrm{SViews}_{s_2}(T_{\Pi_2})$ such that $t_1 \not\sim t_2$ for T_{Π_1} and T_{Π_2} their respective ADL symbolic split-state semantics with respect to some attacker strategy I. We obtain two views witnessing that $\Pi_1\langle s_1\rangle \not\approx_s^{\mathrm{ADL}} \Pi_2\langle s_2\rangle$ from t_1 and t_2 as follows. If there is a consecutive pair of elements of form $((\mathrm{out}, rop, x, y), (\mathrm{in}, x))$ or $((\mathrm{out}, rop, x, y), (\mathrm{in}, \bot))$ with $rop \in \mathcal{ROP}$ and $x, y \in \mathbf{T}$ in t_1 or t_2, then the second element is removed.

Observe that for transition that invoke neither attacker nor library both semantics permit the same transitions. Transitions of the form $(\mathtt{out}, \ldots), (\mathrm{in}, \ldots)$ or $(\mathtt{LC}, \ldots), (\mathrm{libResp}, \ldots)$ correspond immediately, as the library and attacker semantics can be inlined, the only exception being the rule IfTest-m. Wherever IfTest-m is invoked, an event of form $rop \in \mathcal{ROP}$ and $x, y \in \mathbf{T}$ is in the view, hence $V \vdash x$ as well as $V \vdash y$ for V the view at this point. Thus, if this transformation (let us call it θ) is applied, $\theta(t_1) \in \mathrm{SViews}_{s_1}(\Pi_1)$ and $\theta(t_2) \in \mathrm{SViews}_{s_2}(\Pi_2)$. If $t_1 \not\sim t_2$, then $\theta(t_1) \not\sim \theta(t_2)$, as the transformation only removes input elements that are irrelevant for the knowledge set, and is preserving the structure in the following sense: if θ removes the ith element of t_1, it also removes the ith element of t_2, otherwise the $i-1$ prefix of t_1 and t_2 was not symbolically equivalent. □

In Lemma 13, we show that for the embedding presented in Section 7.6 static equivalence in CoSP implies static equivalence in the symbolic variant of ADL.

7.5 Constructing the CoSP-embedding

Previous work [6, 14] defined the embedding into CoSP indirectly, namely via a symbolic and computational execution that followed the structure of a CoSP execution (see Definition 6 and 12). This work, in contrast, explicitly defines an embedding. As CoSP trees are infinite, we define the embedding in a co-recursive manner, i.e., as the largest fixpoint of a co-recursive construction. Each step in this recursion is defined by a function $E_{T_H(s)}$ that takes as input a trace from a leaf-node in the so-far constructed CoSP tree to the root node and outputs a finite subtree. After defining this largest fixpoint construction, we concentrate on defining the recursion $E_{T_H(s)}$. We stress that our construction is defined on the honest program semantics of any over-approximated semantics. Hence, this construction is valid for the canonical symbolic model, too.

7.5.1 Instantiating the over-approximated split-state semantics with references

Within the embedding, we instantiate the honest program semantics in the over-approximated split-state form (see Definition 45), but replace values originating from the cryptographic library or the attacker by pointers to computation nodes or input nodes, respectively. Formally, we instantiate the set $\mathcal{D}_c := Pos$ with the set of positions in a CoSP tree. Here, positions are sequences of natural numbers that encode which path through the CoSP-tree was taken, i.e., $Pos := \mathbb{N}^*$. In the case of ADL, registers and heap locations thus store values that have been input by the adversary or the crypto-API by pointing to the position of the respective input or computation nodes in addition

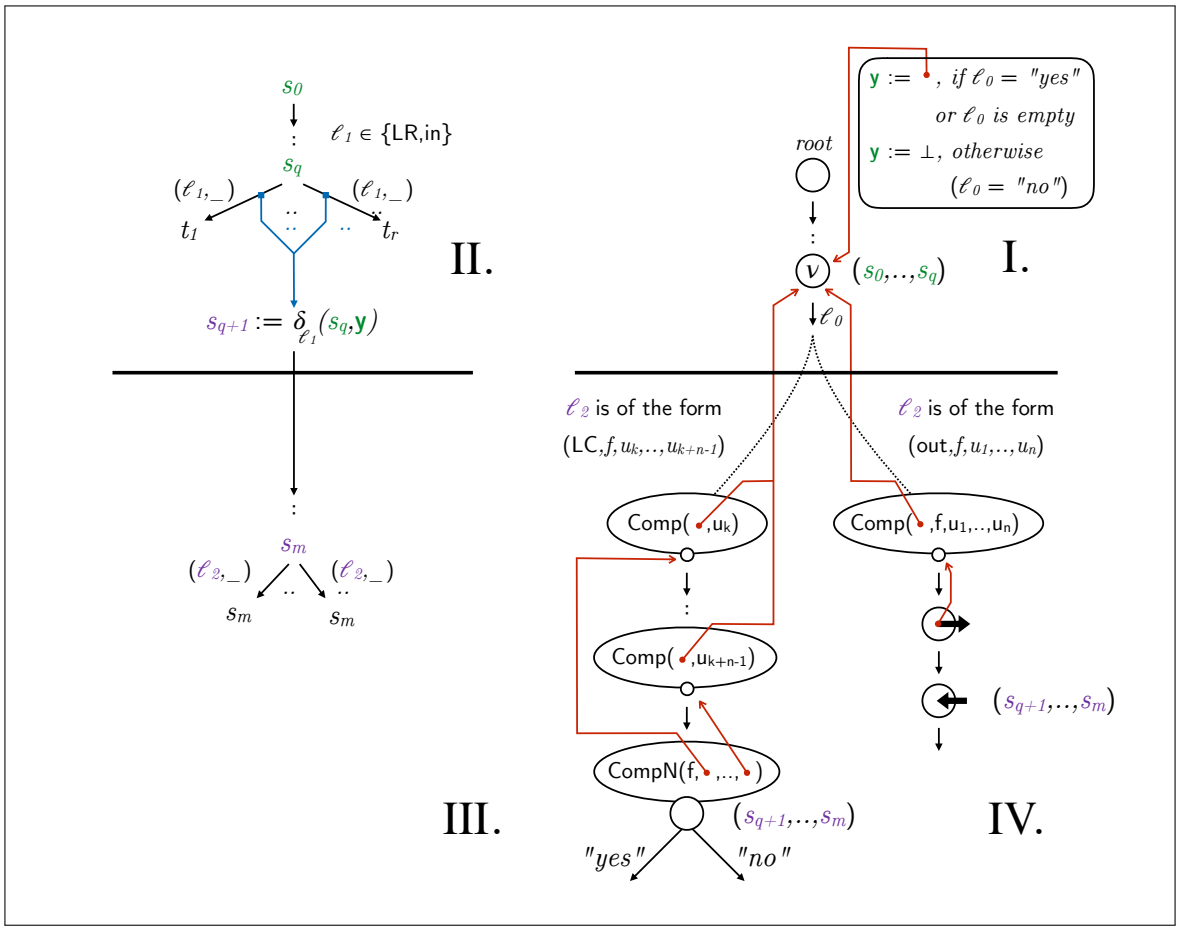

Figure 14: Construction of the recursion step $E_{T_H(s)}$

to numerical values, locations and void. In order to transmit these values to the adversary or the crypto-API, these positions are resolved to a node identifier. The CoSP execution itself takes care of translating, e.g., node identifiers of input nodes to the value the attacker choose to send at this point.

Note that there is no unique representation in $\mathcal{N}'_c$ for all values in $\mathcal{N}$ (i.e., there is no bijection between the sets as required by Definition 3), but (as we will see below) the transitions in Figure 11, which are the only ones using $\mathcal{N}_c$ representations, are never be used. All other transitions are still well-defined.

Recall that states Lemma 4 that the ADL split-state representation is not probabilistic anymore, and contains non-determinism only in global transitions, i.e., transitions with labels in out ∪ in ∪ libCall ∪ libResp.

7.5.2 Constructing the algorithm for the recursion step

The core of the embedding is the recursion step algorithm $E_{T_H(s)}(p, \mathtt{pos})$. As an input, the algorithm expects a path p from a leaf edge ℓ_0 (see Figure 14, phase I.) of the already constructed CoSP tree to the tree's root and the position of the last node in this path. If the leaf node in this already constructed CoSP tree has several outgoing edges, the input p, in particular, also contains one of the outgoing edges (labelled with ℓ_0 in Figure 14). We first give an overview of the recursion step, then describe all auxiliary functions that we use in the full description, and finally present the construction

in Figure 15 in full detail.

Overview. Let the split-state transition system $T = (T_H, T_L, \mathcal{A})$ (i.e., the program in the case of ADL) to be embedded and the initial configuration s be fixed. As Figure 14 depicts, the algorithm can be divided into the following four phases.

I. Extract from the prefix p the label ℓ_0 of the last edge in the CoSP tree that has already been constructed. This last label ℓ_0 determines whether the last subtree that was constructed in the co-recursive computation of the CoSP-tree ended with a computation node or an input node. This check corresponds to split-state transition system ending with a global transition `LR` or `in`. We call this last node ν. We construct $E_{T_H(s)}$ such that ν is annotated with the sequence of states $(s_0, \ldots, s_q)$ of the honest program semantics.[6]
   - If ν is a computation node and $\ell_0 =$ "no", set the variable **y** to the empty reference: **y**$:= \bot$.
   - Otherwise (i.e., $\ell_0 =$ "yes" or ν is an input node), set the variable **y** to the position of ν, as illustrated in Figure 14.

II. Reconstruct a path through the over-approximated honest program semantics T_H (together with the initial state s) with the states $(s_0, \ldots, s_q)$. Check how many possible successors s_q has. $E_{T_H(s)}$ is constructed such that after s_q there is always a global transition, either `LC` or `out`. We determine the successor node s_{q+1} via the global transition mapping δ_{ℓ_1}: $s_{q+1} := \delta_{\ell_1}(s_q, \mathsf{y})$.

III. Run the over-approximated honest program semantics T_H from s_{q+1} until the first global transition ℓ_2.

IV. Construct the corresponding CoSP-subtree as follows.
   - If ℓ_2 is of the form $(\texttt{LC}, f, u_k, ;u_{k+n-1})$, place a sequence C_i of computation nodes that computes the value for each of the arguments u_{k+i} $(i \in 0, \ldots, n-1)$. Then, compute a sequence C_f of computation nodes for the function f that uses the results of the respective sequences $C_0, \ldots, C_{n-1}$. Finally, annotate the final computation nodes of C_f with the sequence $(s_{q+1}, \ldots, s_m)$ of states that occurred in step III.
   - If ℓ_2 is of the form $(\texttt{out}, f, u_1, ;u_n)$, place a sequence C of computation nodes that compute the value for $(f, u_1, \ldots, u_n)$, where f is only an identifier for the malicious function f that triggered the communication to the attacker. Below, place an output node that sends the result of C to the attacker, and one node further place an input node. Annotate the final input node with the sequence $(s_{q+1}, \ldots, s_m)$ of states that occurred in step III.

In the construction of E, we use four auxiliary functions: *exSt*, to extract the honest program semantics-state encoded in a node identifier, *exAllSt*, to extract a list of states that are encoded in a trace of a CoSP tree, t, to find the next global out- or libCall-transition, and id, to find a node for a given position parameter `pos`.

Extracting states: the auxiliary functions *exSt* and *exAllSt*. We define $exSt(\nu)$ as the last state reconstructed out of the identifier ν, and $exAllSt(\nu_1 \to \ldots \nu_n)$ as the function that extract the entire sequence of states from $\nu_1 \to \ldots \nu_n$.

Next global out**- or** libCall**-transition: the auxiliary function t.** We use $p \to$ to denote a path ending at an edge.

$$t(p \to, \hat{s}, n) := \begin{cases} (\hat{s}, n) & \text{if } S(p \to, \hat{s}) = \{s'\} \wedge \hat{s} \xrightarrow{\text{out} \cup \text{libCall}}_H s' \\ t(p \to, s', n+1) & \text{if } S(p \to, \hat{s}) = \{s'\} \text{ but none of the above} \\ (\hat{s}, n) & \text{if } S(p \to, \hat{s}) = \emptyset \end{cases}$$

where S denotes the set of successor states, except when called on the state encoded in the last edge of p, denoted $tl(p)$.

$$S(p \xrightarrow{l}, \hat{s}) := \begin{cases} \{s'\} & \text{if } |\{s'' \mid \hat{s} \to_H s''\}| > 1 \wedge \hat{s} \text{ equals the last state in } exAllSt(p) \wedge \hat{s} \to^l_H s' \\ \{s' \mid \hat{s} \to_H s'\} & \text{otherwise} \end{cases}$$

[6] This sequence of states corresponds to the resulting CoSP subtree from last (co-recursive) invokation of $E_{T_H(s)}$. This overview makes this correspondence precise.

Phase I.
Let $p' := p \to \nu \xrightarrow{[\ell_0]}$. Let $\mathtt{y} := \bot$ if $\ell_0 = \mathtt{no}$ and $\mathtt{y} := id(\mathtt{pos})$ otherwise. Let $(s_0, \ldots, s_q) := exSt(\nu)$.
Phase II.
Let $s_{q+1} := \delta_{\ell_1}(s_q, \mathtt{y})$ if $z = (\ell_1, \_) \wedge \ell_1 \in \{\mathtt{in}, \mathtt{LR}\}$, and $s_0 \to \cdots \to s_q \xrightarrow{z}_H s_{q+1}$
Phase III.
Let ℓ_2 be defined as $t(p', s_{q+1}, 1) = (s_m, m - (q+1)) \wedge s_m \xrightarrow{\ell_2}_H s_m$. In particular, $s_{q+1} \to \cdots \to s_m$.
Phase IV.

$$E_{T_H(s)}(p \to \nu \xrightarrow{[\ell_0]}, \mathtt{pos}) :=$$

$$\begin{cases} c_1 \to^* c_l \to OutN(id(\mathtt{pos}'')) \\ \quad \to InN[((s_{q+1}, \ldots, s_m), \mathtt{pos}')] \to & \text{if } \ell_2 = (\mathtt{out}, f, u_1, \ldots, u_n), \text{ where } \mathtt{pos}' := \mathtt{pos}'' \cdot (0)^2 \\ & \text{and } (c_1 \to^* c_l, \mathtt{pos}'') := Comp(\mathtt{pos}, f, u_1, \ldots, u_n, p') \\ c_1 \to^* c_l \to & \text{if } \ell_2 = (\mathtt{LC}, f, u_k, \ldots, u_{k+n-1}), \text{ where} \\ CompN(f, id(\mathtt{pos}_k), \ldots, id(\mathtt{pos}_{k+n-1})) & (c_1 \to^* c_{l_k}, \mathtt{pos}_k) := Comp(\mathtt{pos}, u_k, p'), \ldots, \\ \quad [((s_{q+1}, \ldots, s_m), \mathtt{pos}')]\left\{\xrightarrow{\mathtt{yes}}, \xrightarrow{\mathtt{no}}\right\} & (c_{1+l_{k+n-2}} \to^* c_l, \mathtt{pos}_{k+n-1}) := Comp(\mathtt{pos}, u_{k+n-1}, p'), \\ & \mathtt{pos}' := \mathtt{pos}_{k+n-1} \cdot (0)^u, \\ & \text{and } u \text{ is the length of } CompN(f, id(\mathtt{pos}_k), \ldots, id(\mathtt{pos}_{k+n-1})) \end{cases}$$

$$Comp(\mathtt{pos}, \underline{u}, p) :=$$

$$\begin{cases} (c_1 \to^* c_k, \mathtt{pos} \cdot (0)^k) & \text{if } \underline{u} = u_1 \wedge\ u_1 \in \mathcal{S} \wedge S \in \{\mathcal{N}, \mathcal{L}, \{\text{void}\}\} \\ & \quad c_1, \ldots, c_k = CoSPRep(\iota_{\mathcal{S}}(u_1), p) \\ ((), u_1) & \text{if } \underline{u} = u_1 \wedge u_1 \in Pos \\ (c_1 \to^* c_k \to^* c_l \to & \text{if } \underline{u} = (u_1, \ldots, u_h) \wedge (c_1 \to^* c_k, \mathtt{pos}_1) = Comp(\mathtt{pos}, u_1, p) \wedge \\ CompN(pair, id(\mathtt{pos}_1, p), id(\mathtt{pos}_2, p)), \mathtt{pos}_2 \cdot (0)) & \quad (c_{k+1} \to^* c_l, \mathtt{pos}_2) = Comp(\mathtt{pos}_1, (u_2, \ldots, u_h), p) \to c_{k+1} \end{cases}$$

$CoSPRep(t, p)$ outputs the representation of a term t in terms of a CoSP path p in an topological sorting of the following tree: for each position in t, there is a node annotated with the constructor at the position, referencing the identifiers of the nodes of its subterms. An edge exists between two nodes iff the term corresponding to the second node is a subterm of the term corresponding to the first node.

$CompN(f, ref_1, \ldots, ref_n)$ is the CoSP tree that corresponds to the symbolic operation f with the references $ref_1, \ldots, ref_n$. Let u be the length of the (linear) CoSP tree that corresponds to f.

Figure 15: Construction of the embedding, where $\nu[(s, \mathtt{pos})]$ means that the identifier of ν is set to $(s, \mathtt{pos})$.

Node for a position: the auxiliary function id. We construct a pair representation for every register name $v \in l$, by first computing the node identifiers for each value at $r(v)$. If $r(v) \in \mathcal{N}_c$, it is a reference to a previous node, which we resolve with the following function, that identifies the position on the current path. We slightly abuse notation by making the position in the node identifier explicit. Note that positions must refer to node on the previous path, as node identifiers, too, can only reference nodes in the prefix of the their path.

$$id(pos, \emptyset) := \bot$$

$$id(pos, p' \to (\hat{s}, pos') \to) := \begin{cases} (\hat{s}, pos) & \text{if } pos = pos' \\ id(pos, p' \to) & \text{otherwise} \end{cases}$$

7.5.3 Corecursively defining the embedding

As outlined above, we are at this point in the position to define the mapping of a (potentially infinite) transition system (the honest program semantics) to an infinite CoSP tree. We define the mapping as the largest fixpoint of a corecursive procedure. For technical reasons (to make the fixpoint unique) and since CoSP trees are necessarily infinite, we first append to the translation of each leaf of the

tree that is spanned by the transition system and an initial configuration an infinite chain of control nodes, which we call a dummy tree.

Definition 52 (Dummy subtree)**.** *A* dummy node *is a control nodes that has as out-metadata its position in the tree and that has a single successor with in-metadata* 0*. A* dummy subtree *is an infinite tree that solely consists of dummy nodes. We call the tree that solely consists of dummy nodes the* empty CoSP protocol*.*

Finally, we can define the embedding of an ADL program as the largest fixpoint as the procedure that step-wise applies Eor appends a dummy subtree if the transition system reaches a final state.

Definition 53 (Embedding of an ADL program)**.** *Let rootpath(edge) be the unique path from the edge to the root. Consider the following definition $U(T_H, p)$:*

- *if $E_{T_H(s)}(p)$ is undefined return a dummy subtree (see Definition 52) to p;*
- *otherwise,*
  - *let $pa := E_{T_H(s)}(p)$,*
  - *append pa to p resulting in a tree par, and*
  - *return the tree where to each leaf le of par the subtree $U(T_H(s), rootpath(le))$ is appended.*

The embedding $U(T_H(s))$ of an honest program transition system T_H with initial configuration s *is co-recursively defined as the greatest fixpoint of $U(T_H(s), \varepsilon)$.*

We stress that this greatest fixpoint $e(T_H(s))$ is unique.

Each CoSP protocol $e(T_H(s))$ is efficient. For each honest program semantics $T_H(s)$ with initial configuration s, $e(T_H(s))$ is efficient in the sense of CoSP. Our algorithm E gives rise to an algorithm that produces an efficient CoSP protocol in the sense of Definition 13, i.e., that only outputs the identifier of the identifier of the next node N and a set of labels for its outgoing edges. This is trivially done by encoding the list of next steps within the identifier.

7.6 Soundness of the CoSP-embedding

We have to show that for every pair of transition system $T_{H,1}, T_{H,2}$ and respective initial configurations s_1, s_2 with $T_{H,1}(s_1) \approx_s^{\mathrm{SS}} T_{H,2}(s_2)$ (i.e., that are equivalent in the split-state semantics) that $T_{H,1}(s_1) \approx_s^{\mathrm{CoSP}} T_{H,2}(s_2)$ holds (i.e., they are is symbolically equivalent in the sense of CoSP). First, we introduce some notation that we use to make the proof more readable, and then we present the soundness proof.

7.6.1 Preliminary definitions

In order to make the proof more readable, we introduce some notation: first the view of an attacker, called out-traces, and second quasi atomic traces, which are used to characterize the subsequences of an honest program semantics execution sequence that is processed by E.

Distribution of out-traces. In the CoSP-computational execution $\langle \mathrm{Exec}_{\mathbf{M},\mathrm{Impl},T_H(s)}(\eta) \| \mathcal{A}(\eta) \rangle$, let $\mathit{OutTraces}_{\mathbf{M},\mathrm{Impl},T_H(s),\mathcal{A},p}(\eta)$ be the distribution of messages that $\mathrm{Exec}_{\mathbf{M},\mathrm{Impl},T_H}(\eta)$ sends to $\mathcal{A}(\eta)$ if output nodes are encountered such that the combined runtime of $\mathrm{Exec}_{\mathbf{M},\mathrm{Impl},T_H}(\eta)$ and $\mathcal{A}(\eta)$ is $\leq p(\eta)$.

Define for transition systems:

- $$\Pr[(a_1, \ldots, a_n) = \mathit{Traces}_{\leq n}(T)] = \sum_{\substack{(\alpha_1, \ldots, \alpha_m)|_A \\ = (a_1, \ldots, a_n) \\ \wedge\, m \leq n}} \Pr[\mathrm{Exec}(T) = s_0 \xrightarrow{\alpha_1}_{m_1} \cdots \xrightarrow{\alpha_m}_{m_m} s_m]$$

- And in general, for random variable X, $\Pr[t = X|_S] := \sum_{t'|_S=t} \Pr[t' = X]$.

Then, let T be any split-state semantics with Crypto-API semantics T_L which harmonizes with the implementation Impl

Let $\mathit{OutTraces}_{SSo,T_H(s),\mathcal{A},p}(\eta)$ be the distribution of out-traces of the transition system $T = (T_H, T_L, \mathcal{A})$ with runtime less or equal than $p(\eta)$, i.e.,

$$\Pr[t = \mathit{OutTraces}_{SSo,T_H,\mathcal{A},s,p}(\eta)] = \sum_{t'|_\pi=t} \Pr[t' = \mathit{Traces}^{ss}_{\eta,z}(T(s))],$$

where $\pi((\text{out}, m)) = m$ and undefined otherwise.

Quasi atomic traces. Given a path ending on an edge, and a state, the first step in the embedding is to iteratively compute the follow-up state until it is no longer uniquely determined, or it involves the attacker or crypto-API. We introduce the notion of an atomic trace and a global transition.

Definition 54 (Global transition)**.** *We call a transition $s \xrightarrow{l} s'$ a* global transition *if $l \in \text{out} \cup \text{in} \cup \text{libCall} \cup \text{libResp}$.*

Definition 55 (Atomic trace)**.** *A sequence $\underline{s}$ of of transition steps $\underline{s} := s_1 \xrightarrow{l}{}^{n} s_n$ is called an* atomic trace *if no global transition occurs. We call a sequence $\underline{s}$ a* quasi atomic trace *if $\underline{s} = s_1 \xrightarrow{l_1} \underline{s}' \xrightarrow{l_2} s_{|\underline{s}'|+2}$ and $l_1 \in \{\texttt{in}, \texttt{LR}\}$ and $l_2 \in \{\texttt{out}, \texttt{LC}\}$.*

7.6.2 The proofs

Lemma 8. *For a sequence $\underline{s}$ of transitions from an honest program semantics T_H that begins at an initial configuration s, let $q_0, \rightarrow_H \cdots \rightarrow_H q_n$ be a partitioning into quasi atomic traces q_i. Let $\rightarrow$ denote an edge in the CoSP tree. Let $e(q_i) := E_{T_H(s)}(q_i)$, and let exAllSt be defined as in Section 7.5.2. With this notation, the following two properties hold.*

(i) For all initial atomic traces q_0, we have $\mathit{exAllSt}(e(q_0)) = q_0$.

(ii) For all $i \in \mathbb{N}$ there is a polynomial p such that we have

$$q_i \rightarrow_H q_{i+1} \implies e(q_i) \rightarrow e(q_{i+1}) \wedge \mathit{exAllSt}(e(q_{i+1})) = q_{i+1}$$

Proof. We show Property (i). Upon each invocation of $E_{T_H(s)}$ a quasi atomic trace or an atomic trace following by one global `out`- or `LC`-transition is generated inside $E_{T_H(s)}$ and stored in the last node; hence Property (i) holds.

Next, we show that Property (ii) also follows from the construction of $E_{T_H(s)}$ and e. If $q_i \rightarrow q_{i+1}$, then e calls $E_{T_H(s)}$ with a prefix $p \rightarrow \nu \xrightarrow{x}$ such that $\mathit{exAllSt}(\nu) = e(q_i)$. Then, $E_{T_H(s)}$ internally computes the transitions beginning from the final state of q_i. The branchings coincide since by Definition 45 $\rightarrow_H$ only branches after `LR`-transitions, which are properly captured by the definition of $E_{T_H(s)}$. We know that the quasi atomic trace beginning from the final state of q_i is a sequence due to the fact that whenever a branching occurs, the library or the adversary is queried. Moreover, the computation in $E_{T_H(s)}$ does not get stuck, despite the fact that the registers carry references instead of values, due to Property (i) of Definition 44. Since $q_i \rightarrow q_{i+1}$, there is a subtree $e(q_i) \rightarrow e(q_{i+1})$ in $e(T_H)$ and as above $\mathit{exAllSt}(e(q_{i+1})) = q_{i+1}$ holds by inspecting the construction of $E_{T_H(s)}$. □

Lemma 9. *Let a symbolic model $\mathbf{M}$ together with an implementation* Impl *be given. For each honest program semantics T_H, for each (attacker) interactive machine $\mathcal{A}$ such that $(T_H, T_L, \mathcal{A})$ satisfies Definition 45, the following holds: there is an interactive machine $R_\mathcal{A}$ and a polynomial p' such that for sufficiently large security parameters η we have*

$$\mathit{OutTraces}_{SSo,T_H,\mathcal{A},s,p}(\eta) = \mathit{OutTraces}_{\mathbf{M},\text{Impl},e(T_H(s)),R_\mathcal{A},p'}(\eta)$$

if $\mathit{OutTraces}_{SSo,T_H,\mathcal{A},p}(\eta)$ uses some Crypto-API $T_L := (T_{L,\eta})_{\eta\in\mathbb{N}}$ that harmonize with Impl *(see Definition 46).*

Proof. We show that for every honest program semantics T_H, every (attacker) interactive machine $\mathcal{A}$, and every resulting out-trace, there is a corresponding out-trace in the CoSP execution. We prove this by induction over the length of the out-trace of the honest program semantics T_H. More formally, for all honest program semantics T_H, we consider a modified transition system $T_H|_n$ that behaves just like T_H except that $T_H|_n$ stops after it produced n out labels. We prove by induction over n that for all honest program semantics T_H and all $n \in \mathbb{N}$, the statement holds for $T_H|_n$; hence the statement holds for polynomial-time computable prefixes.

Constructing the reduction $R_{\mathcal{A}}$. Before we begin with the induction proof, we define an interactive machine $R_{\mathcal{A}}$ from the attacker semantics $\mathcal{A}$.

- The machine $R_{\mathcal{A}}$ internally executes the attacker semantics $\mathcal{A}$ by iteratively computing the successors until either a transition with an `in` label or until there is no unique successor.
- Whenever there is no unique successor but a distribution of successors, $R_{\mathcal{A}}$ conducts a weighted random choice, weighted with the probabilities of each successor.
- Whenever the semantics conducts a transition $s \xrightarrow{(\mathtt{in},m)} s'$, store s' and send m to the communication partner (i.e., the computational CoSP execution).
- Whenever $R_{\mathcal{A}}$ receives a message $f, m_1, \ldots, m_n$ from the interaction partner (i.e., the computational CoSP execution) and the (internal) attacker semantics is currently in state $\langle \varepsilon, as, \varepsilon \rangle$, $R_{\mathcal{A}}$ sets the successor state to be $s' := \langle (f, m_1, \ldots, m_n), as, \varepsilon \rangle$, and using s' the attacker $R_{\mathcal{A}}$ then internally executes the attacker semantics as above.

For notational convenience, we introduce the following notation $\mathrm{outSSo}(i) := \mathit{OutTraces}_{SSo,T_H|_i,\mathcal{A},p}(\eta)$ and $\mathrm{outCOSP}(i) := \mathit{OutTraces}_{\mathbf{M},\mathrm{Impl},e(T_H|_i),R_{\mathcal{A}},p}(\eta)$.

Induction base: $n = 0$. By definition, for all honest program semantics T_H, the support of $\mathrm{outSSo}(0)$ is the empty set. The support of $\mathrm{outCOSP}(0)$ is empty as well, since $E_{T_H|_0(s)}$ (and thus $e(T_H|_0)$) solely produces output nodes whenever there is an `out`-label in the honest program semantics of SSo.

Induction step: $n > 0$. We assume that for all honest program semancs T_H the restricted program $T_H|_{n-1}$ satisfies the statement. Let T_H be an arbitrary but fixed honest program semantics.

Fixing the randomness. It suffices to prove the statement for an arbitrary but fixed set of random choices $r := (r_A, r_L) \in \{0,1\}^*$ in $\mathcal{A}$, written as $\mathcal{A}(r_A)$, and the cryptographic library, written as $\mathrm{Impl}(r_L)$ in CoSP and $SSo(r_L)$ in the over-approximated split-state semantics. We show that for sufficiently large η, for all polynomials p, there is a polynomial p' such that for $r := (r_A, r_L)$

$$\mathrm{outSSo}(n, r) = \mathrm{outCOSP}(n, r),$$

where $\mathrm{outSSo}(i, r) := \mathit{OutTraces}_{\mathrm{ADLo}^{SS}(r_L),T_H|_i,\mathcal{A}(r_A),p}(\eta)$ and $\mathrm{outCOSP}(i, r) := \mathit{OutTraces}_{\mathbf{M},\mathrm{Impl}(r_L),e(T_H|_i),R_{\mathcal{A}(r)},p}(\eta)$.

Reduction to analyzing the last element in $\mathrm{outSSo}(n, r)$. We first show that it suffices to concentrate on proving the statement for the last element in $\mathrm{outSSo}(n, r)$. Recall that $l|_0^k$ denotes the k-prefix of l. If

$$\mathrm{outSSo}(n, r)|_0^{n-1} \neq \mathrm{outCOSP}(n, r)|_0^{n-1}$$

holds, then there is a subset r' of r such that

$$\mathrm{outSSo}(n-1, r') \neq \mathrm{outCOSP}(n-1, r')$$

holds by induction hypothesis.

Analyzing the last element in $\mathrm{outSSo}(n, r)$. For the last element in $\mathrm{outSSo}(n, r)$, we treat all out-transitions uniformly. The last element in $\mathrm{outSSo}(n, r)$ is of the form $(f, ([r(v)]_{\mathcal{N}})_{v \in l})$ for $f \in \mathcal{UNOP} \cup \mathcal{BINOP} \cup \mathcal{RELOP} \cup \mathit{Mal}_{static}$ and $l \in \{(v_b), (v_b, v_c), (v_a, v_b), (v_a, \ldots, v_e)\}$, caused by a honest program

transition $s \xrightarrow{(\mathtt{out}, f, ([r(v)]_{\mathcal{N}})_{v \in l})} s$. $R_{\mathcal{A}(r)}$ behaves just like $\mathcal{A}(r)$ by definition and in particular the input messages are the same. Hence, we know that the inputs, sent to the computational CoSP execution are the same as those that the attacker semantics sends via `in`-transition labels. Moreover, for fixed randomness, the output of the cryptographic library in the CoSP execution are exactly the same as in the library semantics.

As a consequence and by Lemma 8, the state s (from above) is the same (up to the content of the registers) as the state that is internally computed by $E_{T_H|_n(s)}$ before the output node is produced.

Finally, we first stress that the number of computational steps of each invocation of $E_{T_H(s)}$ is polynomially bounded (in η). Hence, for each polynomial p there is a polynomial p' such that the computations coincide. In conclusion, Claim 1 together with definition of $E_{T_H|_n(s)}$ implies that the last element in $\mathrm{out}_{\mathrm{SSo}}(n, r)$ coincides with the last element in $\mathrm{out}_{\mathrm{COSP}}(n, r)$. □

We first prove the preservation in the computational model.

Lemma 10. *Let $T_{H,1}$ and $T_{H,2}$ be two uniform families with respective initial configurations s_1, s_2 of ADL programs compliant (see Definition 47) with the same library specification libSpec w.r.t. to a symbolic model* $\mathbf{M}$ *and an implementation of it, called* Impl. *Let e the injective embedding into CoSP, and s_1 and s_2 initial configurations. Then,*

$$e(T_{H,1}(s_1)) \approx_c^{\mathrm{CoSP}} e(T_{H,2}(s_2)) \implies T_{H,1}(s_1) \approx_c^{\mathrm{SS}} T_{H,2}(s_1).$$

Proof. We assume $e(T_{H,1}(s_1)) \approx_c^{\mathrm{CoSP}} e(T_{H,2}(s_2))$, i.e., for all machines $\mathcal{A}$ and all polynomials p, there is a negligible function μ such that for all $a, b \in \{0,1\}$ with $a \neq b$, all $z \in \{0,1\}^*$ and $\eta \in \mathbb{N}$:

$$\begin{aligned}&\Pr[\langle \mathrm{Exec}_{\mathbf{M},\mathrm{Impl},e(T_{H,1}(s_1))}(\eta) \| \mathcal{A}(\eta, z)\rangle \downarrow_{p(\eta)} a] \\ &+ \Pr[\langle \mathrm{Exec}_{\mathbf{M},\mathrm{Impl},e(T_{H,2}(s_2))}(\eta) \| \mathcal{A}(\eta, z)\rangle \downarrow_{p(\eta)} b] \leq 1 + \mu(\eta).\end{aligned}$$

Proof by contradiction. Assume $T_{H,1}(s_1) \not\approx_c^{\mathrm{SS}} T_{H,2}(s_2)$, i.e., for some family of adversaries $(A^\eta)_{\eta \in \mathbb{N}} \in \mathcal{ADV}$, $\mathrm{ADL}^{SS}_{T^\eta_{H,1}, \mathcal{A}^\eta, s_1}$ and $\mathrm{ADL}^{SS}_{T^\eta_{H,2}, \mathcal{A}^\eta, s_2}$ are distinguishable (see Definition 48). Thus, there are $a, b \in \{0,1\}$, $a \neq b$ and $\eta \in \mathbb{N}$ such that, there is a polynomials p such that for all negligible functions μ,

$$\Pr\left[\mathrm{ADL}^{SS}_{T^\eta_{H,1}, \mathcal{A}^\eta, s_1} \downarrow_{p(\eta)} a\right] + \Pr\left[\mathrm{ADL}^{SS}_{T^\eta_{H,2}, \mathcal{A}^\eta, s_2} \downarrow_{p(\eta)} b\right] > 1 + \mu(\eta).$$

(See Definition 36.) By Corollary 1, there is an attacker A, and a polynomial p' s.t.

$$\Pr\left[\mathrm{ADLo}^{SS}_{T^\eta_{H,1}, \mathcal{A}', s_1} \downarrow_{p(\eta)} a\right] + \Pr\left[\mathrm{ADLo}^{SS}_{T^\eta_{H,2}, \mathcal{A}', s_2} \downarrow_{p(\eta)} b\right] > 1 + \mu(\eta).$$

Thus, by Lemma 9, we obtain the contradiction to the inequality assumed to, concluding the proof. □

In order to match the notation in the overview, we write $e(\Pi_i\langle s_i\rangle)$ for $e(T_{H,i}(s_i))$ in the following lemmas, for two ADL programs Π_1, Π_2 and two initial configurations s_1, s_2.

Lemma 11. *Let Π_1 and Π_2 two ADL program together with input configurations s_1, s_2 in the canonical symbolic model with respect to a CoSP-symbolic model* $\mathbf{M}$*. Let Π_1 and Π_2 be compliant (see Definition 47) with the same library specification w.r.t a symbolic model* $\mathbf{M}$ *and an implementation* Impl. *Then,*

$$e(\Pi_1\langle s_1\rangle) \approx_c^{\mathrm{CoSP}} e(\Pi_2\langle s_2\rangle) \implies \Pi_1\langle s_1\rangle \approx_c^{\mathrm{SS}} \Pi_2\langle s_2\rangle.$$

Next, we can prove the statement for the symbolic side.

Lemma 12. *Let $T_{H,1}$ and $T_{H,2}$ be honest program semantics with respective input configuration s_1, s_2 in the canonical symbolic model with respect to a CoSP-symbolic model* $\mathbf{M}$*. Then,*

$$T_{H,1}(s_1) \approx_s^{\mathrm{SS}} T_{H,2}(s_2) \implies e(T_{H,1}(s_1)) \approx_s^{\mathrm{CoSP}} e(T_{H,2}(s_2))$$

Proof. First, observe that the embedding takes as input any honest program semantics in the over-approximated semantics. Hence, it does exactly the same for the honest program semantics in canonical symbolic model. The same argumentation as in Lemma 9 can be made to show that the states in the canonical symbolic semantics and in the identifiers of the nodes in the embedded program (characterized by the honest program semantics) coincide. The most significant difference is that the attacker is not defined as an interactive machine but rather as a set of deduction rules that are induced by the constructors and the destructors. Instead of fixing an attacker machine, as in the proof of Lemma 9, we fix an attacker strategy (see Definition 6) and then conduct the same induction proof over the length of the deduction sequence (including the steps of the attacker's deductions). Trace equivalence of $T_{H,i}(s_i)$ (for $i \in \{1,2\}$) in the canonical symbolic model and of $e(T_{H,i}(s_i))$ then implies

$$T_{H,1}(s_1) \approx_s^{\mathrm{SS}} T_{H,2}(s_2) \implies e(T_{H,1}(s_1)) \approx_s^{\mathrm{CoSP}} e(T_{H,2}(s_2))$$

□

Lemma 13. *Let Π_1 and Π_2 two ADL program together with input configurations s_1, s_2 in the canonical symbolic model with respect to a CoSP-symbolic model $\mathbf{M}$. Let Π_1 and Π_2 be pre-compliant with the same library specification. Then,*

$$\Pi_1\langle s_1\rangle \approx_s^{\mathrm{SS}} \Pi_2\langle s_2\rangle \implies e(\Pi_1\langle s_1\rangle) \approx_s^{\mathrm{CoSP}} e(\Pi_2\langle s_2\rangle)$$

Proof. The statement immediately follows from Lemma 12. □

Our computational soundness result is parametric in a given symbolic model and given conditions C_I to the implementation and the protocols such that computational soundness in the sense of the CoSP framework holds. Our result states that for any symbolic model with conditions C in CoSP, equivalence in symbolic ADL (see Section 6) implies indistinguishability in ADL (see Section 5). Since all CoSP results in the literature characterize the protocol class by a set of protocol conditions C_P, we use these protocol conditions in our theorem as well.

For our main theorem, we require that two ADL programs are *compliant* (see Definition 47), i.e., use the library in the same way, and the library produces the same distribution as the implementation Impl from the computational soundness result with a library specification *libSpec* (see Definition 46).

Theorem 1. *Let a symbolic model $\mathbf{M}$, protocol conditions C_P and implementation conditions C_I that are computationally sound in the sense of CoSP (Definition 15) be given. Let Π_1 and Π_2 be two uniform families of ADL programs compliant (see Definition 47) with initial configurations s_1, s_2 and the same library specification libSpec w.r.t. to a symbolic model $\mathbf{M}$ and all implementations* Impl *that satisfy C_I. Then, the following implication holds*

$$\Pi_1\langle s_1\rangle \approx_s^{\mathrm{ADL}} \Pi_2\langle s_2\rangle \implies \Pi_1\langle s_1\rangle \approx_c^{\mathrm{ADL}} \Pi_2\langle s_2\rangle.$$

Proof. Follows by transitivity of $\Rightarrow$ from Lemmas 7, 13, 11, and 6, as well as the assumption, that the Impl is a computationally sound implementation of the symbolic model $\mathbf{M}$. □

8 Related work

Operational semantics for Dalvik Bytecode. We have opted to ground our work on the Abstract Dalvik Language (ADL) [4]. ADL currently excels over alternative semantics such as the ones proposed by Wognsen et al. [19], Xia et al. [20], TaintDroid [3], and Chaudhuri [21] because of its comprehensive treatment of the Dalvik language, even though ADL currently only provides sequential executions and does not support exceptions. TaintDroid is more general in this respect in that it contains an ad-hoc modelling of concurrent execution. However, in the ADL extension put forward with the adversary that includes probabilistic choices and adversary interactions, concurrency in ADL can be modelled via program transformations, as we discuss in Section 7.

Information-flow control with cryptographic primitives. The standard notion of security in information flow control – non-interference – is too strong when cryptographic operations are being considered, as, e.g., an encryption of a secret key and a secret message is intuitively safe to be stored in a public variable, but it nontheless results in different values depending on the key and the message. While the declassification of values (see [22] for an introduction and overview of results) can be used to relax this notion, it is difficult to decide under what circumstances cryptographic values can be safely declassified.

Our approach is similar to Askarov et al.'s, permitting so-called *cryptographically masked flows* by considering a relaxed equivalence notion on public values, masking acceptable information flow when ciphertexts are made public [23]. Instead of actually comparing low-values, e.g., encryptions of secret messages, an equivalence relation relates them, e.g., if the encrypted messages have equal length, both encryptions are considered low-equivalent. A security-type system is introduced, which is sound with respect to possibilistic non-interference and said notion of low-equivalence.

Laud has shown the computational soundness of this approach [24] if the employed encryption scheme satisfies key-dependent message security, provides plaintext integrity, and the program is *well-structured*, i.e.,keys are only used in the correct key position, etc. As cryptographically masked flows are captured in a possibilistic setting, leaks through probabilistic behavior are not captured; hence the program must not be able to branch on probabilistic values. Consider the following program:

$$x := \mathsf{rnd}(0,1);\ \mathsf{if}\ x\ \mathsf{then}\ l := h\ \mathsf{else}\ l := \mathsf{rnd}(1,100)$$

While the program is non-interferent in the possibilistic setting, as any final value l has could be random choice in the second branch, it is not non-interferent (insecure) in the probabilistic setting as it is very likely that l indeed contains the value of h.

The consequence for this work is that, since we consider non-interference in the possibilistic setting, we cannot allow any computation on cryptographic values, as these will allow to implement (something close to) rnd using probabilistic operation like encryption. In our work, we treat computations performed on cryptographic values conservatively by relaying them to the adversary. Furthermore, previous work introduced security type-systems, such as [25], and program analyses, such as [26], that directly operate on the computational semantics and are thus capable of verifying non-interference in a probabilistic setting. This approach avoids the aforementioned problem at the cost of less modularity and less potential for automation than our approach.

Computational Soundness. On computational soundness, there is a rich body of literature for various cryptographic primitives [11, 27–29], with malicious keys [8, 30], and even composable computational results [31, 32]. Even though these works covered the applied π-calculus [6, 30, 33], the stateful applied π calculus [34] and RCF [14] and even a embedding of a fragment of C [35], none of these works provide a computational soundness result is known for Dalvik bytecode and hence for Android app analysis.

Another line of work does not provide a complete symbolic characterization of the attacker but concentrates on single rules that hold for certain symbolic terms [36, 37]. This restriction enables a much more flexible and composable preservation notion but pays with omitting any guarantee that the set of rules characterizes all attacker-actions. Hence, it is not clear how well-suited this approach is to automation.

Interactive proof assistants for cryptography. Cryptographic proof assistants enable the mechanized verification of cryptographic proofs, without first abstracting the cryptographic operations [38–40]. Consequently, these tools only offer limited automation. Yet complementarily, these tools could be used to verify that a library satisfies the conditions that a computational soundness result requires.

9 Conclusion and future work

We have shown how cryptographic operations can be faithfully included into existing approaches for automated app analysis. Our results overcome the overly pessimistic results that arise when current automated approaches deal with apps that contain cryptographic operations, as these results do not account for secrecy properties offered by cryptographic operations such as encryption. We have shown how cryptographic operations can be expressed as symbolic abstractions within Dalvik bytecode, so that these abstractions can be conveniently added to existing app analysis tools using minor changes in their semantics. Moreover, we have established the first computational soundness result for the Abstract Dalvik Language (ADL) [4], which currently constitutes the most detailed and comprehensive operational semantics for Dalvik in the literature.

A result that we only scratched on in this work is that any small-step semantics expressed in our novel split-state form entails a canonical small-step semantics for a symbolic model that is computationally sound. This hence provides a recipe for establishing computationally sound symbolic abstractions for *any* given programming language, provided that one can show that the interaction with the attacker and the cryptographic operations can be expressed by means of our concept of split-state semantics. We plan to further investigate this claim and its applicability to modern programming languages in future work.

10 Acknowledgements

This work has been partially funded by the German Research Foundation (DFG) via the collaborative research center "Methods and Tools for Understanding and Controlling Privacy" (SFB 1223), project B3. This work has been partially supported by the Zurich Information Security Center (ZISC).

References

[1] Symantec, "Internet security threat report, volume 20," Accessed: Oct 15, 2015. [Online]. Available: http://www.symantec.com/security_response/publications/threatreport.jsp

[2] GData, "Mobile malware report: Q2/2015," Accessed: Oct 15, 2015. [Online]. Available: https://public.gdatasoftware.com/Presse/Publikationen/Malware_Reports/G_DATA_MobileMWR_Q2_2015_EN.pdf

[3] W. Enck, P. Gilbert, B.-G. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth, "Taintdroid: an information flow tracking system for real-time privacy monitoring on smartphones," *Communications of the ACM*, vol. 57, no. 3, pp. 99–106, 2014.

[4] S. Lortz, H. Mantel, A. Starostin, T. Bähr, D. Schneider, and A. Weber, "Cassandra: towards a certifying app store for android," in *Proc. 4th ACM SPSM*, 2014, pp. 93–104.

[5] M. Backes, S. Bugiel, E. Derr, S. Gerling, and C. Hammer, "R-droid: In-depth application vetting for android with path-sensitive value analysis," in *Proc. ACM AsiaCCS*, 2016.

[6] M. Backes, D. Hofheinz, and D. Unruh, "CoSP: A General Framework for Computational Soundness Proofs," in *Proc. 16th ACM CCS*, 2009, pp. 66–78.

[7] S. Meiser, "Computational soundness of passively secure encryption in presence of active adversaries," Master's thesis at Saarland University, 2010.

[8] M. Backes, A. Malik, and D. Unruh, "Computational Soundness without Protocol Restrictions," in *Proc. 19th ACM CCS*, 2012, pp. 699–711.

[9] M. Backes, F. Bendun, and D. Unruh, "Computational Soundness of Symbolic Zero-knowledge Proofs: Weaker Assumptions and Mechanized Verification," in *Proc. 2nd POST*, 2013, pp. 206–225.

[10] M. Backes, E. Mohammadi, and T. Ruffing, "Computational Soundness Results for ProVerif," in *Proc. 3rd POST*, 2014, pp. 42–62.

[11] M. Backes, F. Bendun, M. Maffei, E. Mohammadi, and K. Pecina, "A Computationally Sound, Symbolic Abstraction for Malleable Zero-knowledge Proofs," in *Proc. 28th IEEE CSF*, 2015, pp. 412–480.

[12] D. Unruh, "Termination-Insensitive Computational Indistinguishability (and Applications to Computational Soundness)," in *Proc. 24th IEEE CSF*, 2011, pp. 251–265.

[13] M. Y. Vardi, "Automatic verification of probabilistic concurrent finite state programs," in *Proc. 26th SFCS*, 1985, pp. 327–338.

[14] M. Backes, M. Maffei, and D. Unruh, "Computationally Sound Verification of Source Code," in *Proc. 17th ACM CCS*, 2010, pp. 387–398.

[15] M. Aizatulin, A. D. Gordon, and J. Jürjens, "Extracting and verifying cryptographic models from C protocol code by symbolic execution," in *Proc. 18th ACM CCS*, 2011, pp. 331–340.

[16] M. Backes, R. Künnemann, and E. Mohammadi, "Technical report: Computational soundness for dalvik bytecode," arXiv:1608.04362, 2016.

[17] R. Milner, *Communication and concurrency*, 1989, vol. 84.

[18] C. A. R. Hoare *et al.*, *Communicating sequential processes*, 1985, vol. 178.

[19] E. R. Wognsen, H. S. Karlsen, M. C. Olesen, and R. R. Hansen, "Formalisation and analysis of dalvik bytecode," *Science of Computer Programming*, vol. 92, pp. 25–55, 2014.

[20] M. Xia, L. Gong, Y. Lyu, Z. Qi, and X. Liu, "Effective real-time android application auditing," in *Proc. 36th IEEE S&*, 2015, pp. 899–914.

[21] A. Chaudhuri, "Language-based security on android," in *Proc. 4th ACM PLAS*, 2009, pp. 1–7.

[22] A. Sabelfeld and D. Sands, "Declassification: Dimensions and principles," *J. Computer Security*, vol. 17, no. 5, 2009.

[23] A. Askarov, D. Hedin, and A. Sabelfeld, "Cryptographically-masked flows," *TCS*, vol. 402, no. 2-3, pp. 82–101, 2008.

[24] P. Laud, "On the computational soundness of cryptographically masked flows," in *Proc. of 35th POPL*, 2008, pp. 337–348.

[25] P. Laud and V. Vene, "A type system for computationally secure information flow," in *Proc. of 15th FCT*, 2005, pp. 365–377.

[26] P. Laud, "Handling encryption in an analysis for secure information flow," in *Proc. of 12th ESOP*, 2003, pp. 159–173.

[27] M. Backes, B. Pfitzmann, and M. Waidner, "A Composable Cryptographic Library with Nested Operations," in *Proc. 10th ACM CCS*, 2003, pp. 220–230.

[28] V. Cortier and B. Warinschi, "Computationally Sound, Automated Proofs for Security Protocols," in *Proc. 14th ESOP*, 2005, pp. 157–171.

[29] D. Galindo, F. D. Garcia, and P. van Rossum, "Computational Soundness of Non-Malleable Commitments," in *Proc. 4th ISPEC*, 2008, pp. 361–376.

[30] H. Comon-Lundh, V. Cortier, and G. Scerri, "Security Proof with Dishonest Keys," in *Proc. 1nd POST*. Springer, 2012, pp. 149–168.

[31] V. Cortier and B. Warinschi, "A composable computational soundness notion," in *Proc. 18th ACM CCS*, 2011, pp. 63–74.

[32] F. Böhl, V. Cortier, and B. Warinschi, "Deduction Soundness: Prove One, Get Five for Free," in *Proc. 20th ACM CCS*, 2013, pp. 1261–1272.

[33] H. Comon-Lundh and V. Cortier, "Computational Soundness of Observational Equivalence," in *Proc. 15th ACM CCS*. ACM Press, 2008, pp. 109–118.

[34] J. Shao, Y. Qin, and D. Feng, "Computational Soundness Results for Stateful Applied Pi Calculus," in *Proc. 5rd POS*, 2016, pp. 254–275.

[35] M. Aizatulin, A. D. Gordon, and J. Jürjens, "Computational Verification of C Protocol Implementations by Symbolic Execution," in *Proc. 19th ACM CCS*, 2012, pp. 712–723.

[36] G. Bana and H. Comon-Lundh, "Towards Unconditional Soundness: Computationally Complete Symbolic Attacker," in *Proc. 1nd POST*, 2012, pp. 189–208.

[37] ——, "A Computationally Complete Symbolic Attacker for Equivalence Properties," in *Proc. 21th ACM CCS*, 2014, pp. 609–620.

[38] G. Barthe, B. Grégoire, S. Heraud, and S.-Z. Béguelin, "Computer-Aided Security Proofs for the Working Cryptographer," in *Proc. CRYPTO*, 2011, pp. 71–90.

[39] A. Petcher and G. Morrisett, "The Foundational Cryptography Framework," in *Proc. 4th POST*, 2015, pp. 53–72.

[40] A. Lochbihler, "Probabilistic functions and cryptographic oracles in higher order logic," in *Proc. 25th ESOP*, 2016, pp. 503–531.

$$\text{rNop}\ \frac{m[pp] = \mathsf{nop}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h, pp+1, r\rangle}$$

$$\text{rGoto}\ \frac{m[pp] = \mathsf{goto}\ n}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h, pp+n, r\rangle}$$

Figure 16: Semantics of control flow instructions

Appendix

A Inference rules for the DEX Bytecode Semantics

B Extendability

Definition 56 (Efficient transparent function)**.** *An n-ary constructor f is* transparent *if for every argument $i \in \{1, \ldots, n\}$ there is a destructor $d_{inverse,i}$, called* the i-th inverse function of f*, such that for all terms $M_1, \ldots, M_n$, we have $eval_{d_{inverse,i}} f(M_1, \ldots, M_n) = M_i$.*

Definition 57 (Finitely generated message types)**.** *A message type $\mathbf{T}$ is* finitely generated *if there exists a grammar with rules $r_1, \ldots, r_n$ such that $\mathbf{T}$ equals the set that is generated by $r_1, \ldots, r_n$.*[7]

Definition 58 (Combination of two finitely generated message types)**.** *Let $\mathbf{T}$ be a finitely generated message types over $\mathbf{C}$ and $\mathbf{N}$ with a grammar $r_1, \ldots, r_n$, and let $\mathbf{T}'$ be a finitely generated message types over $\mathbf{C}'$ and $\mathbf{N}'$ with a grammar $r'_1, \ldots, r'_m$ such that $\mathbf{C} \cap \mathbf{C}' = \emptyset$ and $\mathbf{N} \cap \mathbf{N}' = \emptyset$. Then, the* combination $\mathbf{T} \sqcup \mathbf{T}'$ of $\mathbf{T}$ and $\mathbf{T}'$ *is the set that is generated by the rules $r_1, \ldots, r_n, r'_1, \ldots, r'_m$.*

Definition 59 (Combination of two symbolic models)**.** *Let a symbolic model $\mathbf{M} = (\mathbf{C}, \mathbf{N_E} \uplus \mathbf{N_P}, \mathbf{T}, \mathbf{D})$ and another symbolic model $\mathbf{M}' = (\mathbf{C}', \mathbf{N_E}' \uplus \mathbf{N_P}', \mathbf{T}', \mathbf{D}')$ with $\mathbf{C} \cap \mathbf{C}' = \emptyset$, $\mathbf{D} \cap \mathbf{D}' = \emptyset$, $\mathbf{N_E} \cap \mathbf{N_E}' = \emptyset$, $\mathbf{N_P} \cap \mathbf{N_P}' = \emptyset$, $\mathbf{T} \cap \mathbf{T}' = \emptyset$ be given. Then, the* combination of $\mathbf{M}$ and $\mathbf{M}'$ *is defined as $\mathbf{M} \sqcup \mathbf{M}' := (\mathbf{C} \cup \mathbf{C}', \mathbf{N} = \mathbf{N_E} \cup \mathbf{N_E}' \uplus \mathbf{N_P} \cup \mathbf{N_P}', \mathbf{T} \sqcup \mathbf{T}', \mathbf{D} \cup \mathbf{D}')$.*

Definition 60 (Restriction-defined protocol class)**.** *Let a symbolic model $\mathbf{M} = (\mathbf{C}, \mathbf{N}, \mathbf{T}, \mathbf{D})$ be given. We call a set $R := \{c_1, \ldots, c_n\}$ of functions from CoSP protocols to $\{\top, \bot\}$, i.e., true or false,* protocol restrictions*. A protocol class* P *is* restriction-defined *by R on $\mathbf{M}$, written as* $\mathrm{P} := \mathcal{P}(R, \mathbf{M})$*, if* P *is the largest set of protocols for $\mathbf{M}$ (see Definition 4) such that $\forall \Pi \in \mathrm{P}. \forall i \in \{1, \ldots, n\}. c_i(\Pi) = \top$.*

Lemma 14. *Given a symbolic model $\mathbf{M}$ with a implementation* Impl *of $\mathbf{M}$. Let $\mathbf{M}_{trp} = (\mathbf{C}', \mathbf{N}', \mathbf{T}', \mathbf{D}')$ be a symbolic model with an implementation Impl_{trp} such that all constructors $f \in \mathbf{C}'$ are transparent functions. Let R be protocol restrictions (see Definition 60).*

If the implementation Impl *is computationally sound for $\mathbf{M}$ and the protocol class $\mathcal{P}(R, \mathbf{M})$, then the implementation $\mathrm{Impl} \cup \mathrm{Impl}_{trp}$ is computationally sound for the symbolic model $\mathbf{M} \sqcup \mathbf{M}_{trp}$ and the protocol class $\mathcal{P}(R, \mathbf{M} \sqcup \mathbf{M}_{trp})$.*

Proof. By contraposition it suffices to show that for any pair of protocols $\Pi'_1, \Pi'_2 \in \mathcal{P}(R, \mathbf{M} \sqcup \mathbf{M}_{trp})$, the following holds

$$\Pi'_1 \not\approx_c^{\mathrm{CoSP}} \Pi'_2 \implies \Pi'_1 \not\approx_s^{\mathrm{CoSP}} \Pi'_2$$

Towards a contradiction, assume that

$$\Pi'_1 \not\approx_c^{\mathrm{CoSP}} \Pi'_2 \wedge \Pi'_1 \approx_s^{\mathrm{CoSP}} \Pi'_2$$

holds. Our strategy is to show that there are protocols $\Pi_1, \Pi_2 \in \mathcal{P}(R, \mathbf{M})$ such that

[7] The notion of a set that is generated by a grammar is a common notion. It refers to the smallest fixpoint that obeys the rules from the grammar.

$$\text{rInstanceOfTrue}\ \frac{\begin{array}{c} m[pp] = \textsf{instance-of}\ v_a, v_b, cl \qquad r(v_b) \in dom(h) \\ assignmentCompatible(h(r(v_b))).cls, cl \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h, pp+1, r[v_a \mapsto 1]\rangle}$$

$$\text{rInstanceOfFalse}\ \frac{\begin{array}{c} m[pp] = \textsf{instance-of}\ v_a, v_b, cl \qquad r(v_b) \in dom(h) \\ \neg(assignmentCompatible(h(r(v_b)).cls, cl)) \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \mapsto \langle h, pp+1, r[v_a \mapsto 0]\rangle}$$

$$\text{rNewInstance}\ \frac{\begin{array}{c} m[pp] = \textsf{new-instance}\ v_a, cl \qquad h \in dom(nextFreeLocation) \\ l = nextFreeLocation(h) \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h[l \mapsto defaultObject(cl)], pp+1, r[v_a \mapsto l]\rangle}$$

$$\text{rConstString}\ \frac{m[pp] = \textsf{const-string}\ v_a, s \qquad s \in dom(nameToReference)}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h, pp+1, r[v_a \mapsto nameToReference(s)]\rangle}$$

$$\text{rConstClass}\ \frac{m[pp] = \textsf{const-class}\ v_a, cl \qquad cl \in dom(nameToReference)}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h, pp+1, r[v_a \mapsto nameToReference(cl)]\rangle}$$

$$\text{rIget}\ \frac{\begin{array}{c} m[pp] = \textsf{iget}\ v_a, v_b, fid \qquad fid \in dom(lookup\text{-}field_\Pi) \\ r(v_b) \in dom(h) \qquad o = h(r(v_b)) \\ f = lookup\text{-}field_\Pi(fid) \qquad f \in dom(o.fields) \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h, pp+1, r[v_a \mapsto o.f]\rangle}$$

$$\text{rIput}\ \frac{\begin{array}{c} m[pp] = \textsf{iput}\ v_a, v_b, fid \qquad fid \in dom(lookup\text{-}field_\Pi) \\ r(v_b) \in dom(h) \qquad o = h(r(v_b)) \\ f = lookup\text{-}field_\Pi(fid) \qquad f \in dom(o.fields) \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h[r(v_b) \mapsto o[f \mapsto r(v_a)]], pp+1, r\rangle}$$

$$\text{rSget}\ \frac{\begin{array}{c} m[pp] = \textsf{sget}\ v_a, fid \qquad fid \in dom(nameToReference) \\ l = nameToReference(fid) \qquad fid \in dom(lookup\text{-}field_\Pi) \\ f = lookup\text{-}field_\Pi(fid) \qquad f \in dom(h(l).fields) \qquad u = h(l).f \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h, pp+1, r[v_a \mapsto u]\rangle}$$

$$\text{rSput}\ \frac{\begin{array}{c} m[pp] = \textsf{sput}\ v_a, fid \qquad fid \in dom(nameToReference) \\ l = nameToRefrence(fid) \qquad fid \in dom(lookup\text{-}field_\Pi) \\ f = lookup\text{-}field_\Pi(fid) \qquad o = h(l) \qquad f \in dom(o.fields) \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h[l \mapsto o[f \mapsto r(v_a)]], pp+1, r\rangle}$$

Figure 17: Semantics of object-related instructions

$$\Pi_1' \not\approx_c^{\text{CoSP}} \Pi_2' \implies \Pi_1 \not\approx_c^{\text{CoSP}} \Pi_2 \text{ (Claim 1)}$$
$$\Pi_1 \not\approx_s^{\text{CoSP}} \Pi_2 \implies \Pi_1' \not\approx_s^{\text{CoSP}} \Pi_2' \text{ (Claim 2)}$$

holds. These two statements imply that

$$\Pi_1 \not\approx_c^{\text{CoSP}} \Pi_2 \wedge \Pi_1 \approx_s^{\text{CoSP}} \Pi_2$$

holds, which contradicts the computational soundness of Impl for $\mathbf{M}$ and $\mathcal{P}(R, \mathbf{M})$.

Claim 1. $\Pi_1' \not\approx_c^{\text{CoSP}} \Pi_2' \implies \Pi_1 \not\approx_c^{\text{CoSP}} \Pi_2$

Proof of Claim 1. First, we define the *encoding of an n-ary transparent constructor* f as follows:

$$encoding - f(\nu_1, \ldots, \nu_n) := pair(\texttt{constructor-f}, pair(\nu_1, pair(, \ldots, pair(\nu_{n-1}, \nu_n))))$$

Analogously, we define the *encoding of the i-th inverse function* $d_{inverse,i}$ *of* f as the sequence of the *fst* and *snd* destructor that projects to the i-th argument. We denote this encoding as $encoding - d - inverse - i(\nu_1)$.

$$\text{rArrayLength}\frac{m[pp] = \textsf{array-length}\ v_a, v_b \quad r(v_b) \in dom(h)}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h, pp+1, r[v_a \mapsto h(r(v_b)).length]\rangle}$$

$$\text{rNewArray}\frac{\begin{array}{c} m[pp] = \textsf{new-array}\ v_a, v_b \quad h \in dom(nextFreeLocation) \\ l = nextFreeLocation(h) \qquad 0 \leq r(v_b) \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h[l \mapsto defaultArray(r(v_b))], pp+1, r[v_a \mapsto l]\rangle}$$

$$\text{rFilledNewArrayR}\frac{\begin{array}{c} m[pp] = \textsf{filled-new-array-range}\ v_k, n \quad h \in dom(nextFreeLocation) \\ l = nextFreeLocation(h) \qquad x = defaultArray(n) \\ ar = x[0 \mapsto r(v_k), \ldots, n-1 \mapsto r(v_{k+n-1})] \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h[l \mapsto ar], pp+1, r[res_{lo} \mapsto l]\rangle}$$

$$\text{rAget}\frac{\begin{array}{c} m[pp] = \textsf{aget}\ v_a, v_b, v_c \quad r(v_b) \in dom(h) \quad ar = h(r(v_b)) \\ u = ar[r(v_c)] \qquad 0 \leq r(v_c) < ar.length \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h, pp+1, r[v_a \mapsto u]\rangle}$$

$$\text{rAput}\frac{\begin{array}{c} m[pp] = \textsf{aput}\ v_a, v_b, v_c \quad r(v_b) \in dom(h) \quad ar = h(r(v_b)) \\ x = ar[r(v_c) \mapsto r(v_a)] \qquad 0 \leq r(v_c) < ar.length \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \langle h[r(v_b) \mapsto x], pp+1, r\rangle}$$

Figure 18: Semantics of array-related instructions

$$\text{rIVR}\frac{\begin{array}{c} m[pp] = \textsf{invoke-virtual-range}\ v_k, n, mid \qquad r(v_k) \in dom(h) \\ (mid, h(r(v_k)).cls) \in dom(lookup\text{-}virtual_\Pi) \quad m' = lookup\text{-}virtual_\Pi(mid, h(r(v_k)).cls) \\ \langle h, 0, defaultRegisters([r(v_k), \ldots, r(v_{k+n-1})])\rangle \Downarrow^{(n')}_{P,m'} \langle u, h'\rangle \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow_{n'+1} Pm\, \langle h', pp+1, r[res_{lo} \mapsto lower(u), res_{up} \mapsto upper(u)]\rangle}$$

$$\text{rIStR}\frac{\begin{array}{c} m[pp] = \textsf{invoke-static-range}\ v_k, n, mid \\ mid \in dom(lookup\text{-}static_\Pi) \qquad m' = lookup\text{-}static_\Pi(mid) \\ \langle h, 0, defaultRegisters([r(v_k), \ldots, r(v_{k+n-1})])\rangle \Downarrow^{(n')}_{P,m'} \langle u, h'\rangle \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow_{n'+1} Pm\, \langle h', pp+1, r[res_{lo} \mapsto lower(u), res_{up} \mapsto upper(u)]\rangle}$$

Figure 19: Semantics of method-related instructions

We inductively define Π_i to equal Π'_i except for occurrences of computation nodes ν that are labelled with a transparent constructor or an inverse function, which are in Π_i replaced by the sequence of computation nodes that represents the encoding of the transparent constructor or its inverse function.

The claim follows since for all terms $t_1, \ldots, t_n$ the term $encoding - d - inverse - i(encoding - f(t_1, \ldots, t_n))$ evaluates to t_i. ◇

Claim 2. $\Pi_1 \not\approx_s^{\text{CoSP}} \Pi_2 \implies \Pi'_1 \not\approx_s^{\text{CoSP}} \Pi'_2$

Proof of Claim 2. Analogous to the proof of Claim 1, Claim 2 follows since for all terms $t_1, \ldots, t_n$ the term $encoding - d - inverse - i(encoding - f(t_1, \ldots, t_n))$ evaluates to t_i. ◇

□

B.1 Arbitrary bit-operations on bitstrings

We can extend any computationally sound symbolic model that contains the basic bitstring operations $string_0/1, string_1/1, unstring_0/1, unstring_1/1, emp/0$ with

$$\forall b \in \{0,1\}. unstring_b(string_b(x)) = x$$

to a symbolic model that contains all bitstring operations $op : \{0,1\}^* \rightarrow \{0,1\}^*$ as long as the respective destructor is solely defined on symbolic string. The proof encodes any operation on bitstring

$$\text{rUnopWideS}\ \frac{\begin{array}{c} m[pp] = \mathsf{unop\text{-}wideS}\ v_a, v_b, uop \\ u = \underline{uop}(r(v_b) \bullet r(v_{b+1})) \end{array}}{\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow\ \langle h, pp+1, r[v_a \mapsto u]\rangle}$$

$$\text{rUnopWideT}\ \frac{\begin{array}{c} m[pp] = \mathsf{unop\text{-}wideT}\ v_a, v_b, uop \\ u = \underline{uop}(r(v_b)) \end{array}}{\begin{array}{l}\langle m \cdot ml, h, pp \cdot ppl, r \cdot rl, as\rangle \rightarrow \\ \langle h, pp+1, r[v_a \mapsto lower(u), v_{a+1} \mapsto upper(u)]\rangle\end{array}}$$

Figure 20: Semantics of conversion instructions for 64 bit values

to the corresponding circuit. Hence, we will restrict ourselves to illustrating how $\wedge, \vee, \neg$ gates are encoded on single bits:

$\wedge(string_{b_1}(emp), string_{b_2}(emp))$

1: **if** $string_{b_1}(emp) = string_0(emp) \wedge string_{b_2}(emp) = string_0(emp)$ **then**
2: **return** $string_0(emp)$
3: **else if** $string_{b_1}(emp) = string_1(emp) \wedge string_{b_2}(emp) = string_1(emp)$ **then**
4: **return** $string_1(emp)$
5: **else if** $string_{b_1}(emp) = string_0(emp) \wedge string_{b_2}(emp) = string_1(emp)$ **then**
6: **return** $string_0(emp)$
7: **else if** $string_{b_1}(emp) = string_1(emp) \wedge string_{b_2}(emp) = string_0(emp)$ **then**
8: **return** $string_0(emp)$

$\vee(string_{b_1}(emp), string_{b_2}(emp))$

1: **if** $string_{b_1}(emp) = string_0(emp) \wedge string_{b_2}(emp) = string_0(emp)$ **then**
2: **return** $string_0(emp)$
3: **else if** $string_{b_1}(emp) = string_1(emp) \wedge string_{b_2}(emp) = string_1(emp)$ **then**
4: **return** $string_1(emp)$
5: **else if** $string_{b_1}(emp) = string_0(emp) \wedge string_{b_2}(emp) = string_1(emp)$ **then**
6: **return** $string_1(emp)$
7: **else if** $string_{b_1}(emp) = string_1(emp) \wedge string_{b_2}(emp) = string_0(emp)$ **then**
8: **return** $string_1(emp)$

$\neg(string_b(emp))$

1: **if** $string_b(emp) = string_0(emp)$ **then**
2: **return** $string_1(emp)$
3: **else if** $string_b(emp) = string_1(emp)$ **then**
4: **return** $string_0(emp)$

These encodings can be easily extended to bitstrings and used in order to encode arbitrary circuits, hence any computable function $op : \{0,1\}^* \rightarrow \{0,1\}^*$, on bitstrings. Given a bijection ι from bitstrings to symbolic bitstrings as above, a symbolic binary operation d_{op}/n is a destructor that is solely defined on symbolic bitstrings as

$$d_{op}(x_1, \ldots, x_n) = \iota(op(\iota^{-1}(x_1), \ldots, \iota^{-1}(x_n)))$$

As a corollary, we get

Corollary 2. *Given a symbolic model* $\mathbf{M}$ *with a implementation* Impl *of* $\mathbf{M}$*. Let* $\mathbf{M}_{bitstringOp} = (\emptyset, \mathbf{N}', \mathbf{T}', \mathbf{D}')$ *be a symbolic model with an implementation* $\text{Impl}_{bitstringOp}$ *such that all destructors* $f \in \mathbf{D}'$ *are binary operations. Let* R *be protocol restrictions (see Definition 60).*

If the implementation Impl *is computationally sound for* $\mathbf{M}$ *and the protocol class* $\mathcal{P}(R, \mathbf{M})$*, then the implementation* $\text{Impl} \cup \text{Impl}_{bitstringOp}$ *is computationally sound for the symbolic model* $\mathbf{M} \sqcup \mathbf{M}_{bitstringOp}$ *and the protocol class* $\mathcal{P}(R, \mathbf{M} \sqcup \mathbf{M}_{trp})$*.*

B.2 Derived destructors with symbolic operations

Similar to transparent functions and bitstring operations on symbolic bitstrings, all destructors that can be represented as a symbolic operation can be added to a computationally sound symbolic model as destructors applied to nonces, i.e., destructor and nonce computation nodes, without loosing computational soundness. With the formal parameters as holes, a symbolic operation O, with say n variables, can be interpreted as a context over terms. Such a context naturally defines an n-ary destructor that is applied to nonces $r_{k+1}, \ldots, r_n$ via $eval_O(t_1, \ldots, t_k, r_{k+1}, \ldots, r_n)$.[8] We call such a destructor a *derived destructor*.

Such derived destructors can be encoded using constructors and destructors from the original symbolic model, we can use the same proof techniques as above and conclude computational soundness for the combined model.

Corollary 3. *Given a symbolic model* $\mathbf{M}$ *with a implementation* Impl *of* $\mathbf{M}$. *Let* $\mathbf{M}_{derDes} = (\emptyset, \mathbf{N}', \mathbf{T}', \mathbf{D}')$ *be a symbolic model with an implementation* Impl_{derDes} *such that all destructors* $f \in \mathbf{D}'$ *are derived destructors and the implementation is the evaluation of the symbolic operation, using the implementations of the constructors and destructors that occur in the symbolic operation. Let* R *be protocol restrictions (see Definition 60).*

If the implementation Impl *is computationally sound for* $\mathbf{M}$ *and the protocol class* $\mathcal{P}(R, \mathbf{M})$*, then the implementation* $\mathrm{Impl} \cup \mathrm{Impl}_{derDes}$ *is computationally sound for the symbolic model* $\mathbf{M} \sqcup \mathbf{M}_{derDes}$ *and the protocol class* $\mathcal{P}(R, \mathbf{M} \sqcup \mathbf{M}_{trp})$.

[8] Without loss of generality, we assumed here an order on the parameters.